Project Status Report

High End Computing Capability
Strategic Capabilities Assets Program

September 10, 2015

Dr. Rupak Biswas – Project Manager
NASA Ames Research Center, Moffett Field, CA
Rupak.Biswas@nasa.gov
(650) 604-4411

Pleiades Surpasses a Half Billion SBUs Delivered to Users

- NASA's flagship supercomputer, Pleiades, delivered 16,535,458 Standard Billing Units (SBUs) in August, bringing the Pleiades total to 502,620,650 SBUs.
- HECC deployed multiple expansions to Pleiades, increasing its computational capability by about 8x since its installation in 2008. Pleiades now delivers in 6 weeks the same amount of processing that its predecessor, Columbia, delivered in its entire 9-year life span, which ended in 2013.
- Pleiades already delivered ~25x the processing Columbia delivered. Columbia delivered approximately 10x the combined processing of all NAS supercomputers before it, going back to the establishment of the NAS facility in 1987.
- Among recent projects using the most computer time on Pleiades: flow computations for the Space Launch System configuration; validation of Kepler mission planet candidates; investigations to reduce the environmental impact of aircraft; and global simulations to help monitor ocean, sea ice, and atmospheric systems.

** 1 SBU equals 1 hour of a Pleiades Westmere 12-core node.*

Mission Impact: The HECC project's continuous increase in computational capability and capacity is essential for all mission directorates to meet the high demand for supercomputing resources to accomplish their goals and objectives.

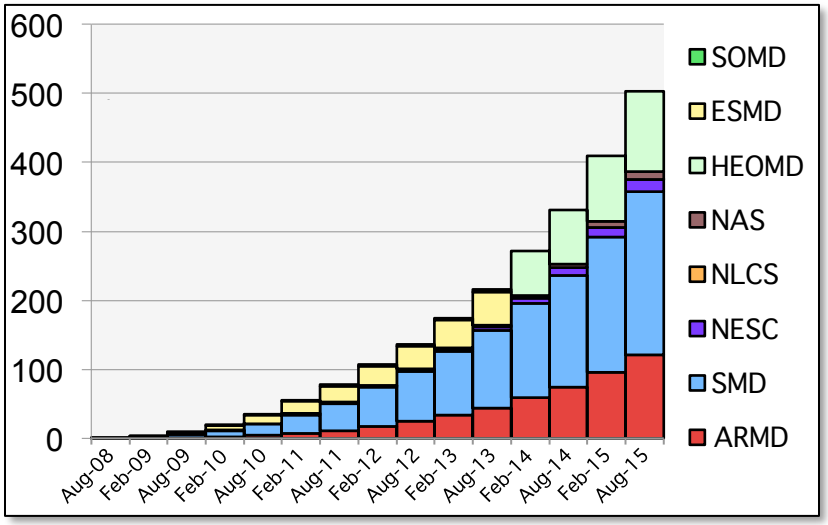

Cumulative chart in millions since the first mission directorate jobs ran on Pleiades from August 2008 to August 2015. One hour of computer time is defined as one hour of computing on a 12-core Westmere node. Therefore, the 500-million hour milestone is equivalent to delivering 6 billion Westmere core hours.

POC: Catherine Schulbach, catherine.h.schulbach@nasa.gov, (650) 604-3180, NASA Advanced Supercomputing Division

Facilities Team Coordinates Successful Power Maintenance Shutdown

- HECC staff planned and executed a shutdown of the secondary compute facility (building N233A) at Ames Research Center (ARC). During the shutdown, ARC upgraded the relay switches on four power transformers.
- N233A houses critical HECC computing resources, including the Merope supercomputer, tape silos, disk storage, and security servers.
- Facilities worked with contractors performing the work, and all affected support groups (Systems, Security, Engineering Servers and Services, Networks, NASA Integrated Communication Services) to orchestrate the smooth shutdown and startup of these HECC resources.
- Following a comprehensive planning/scheduling process, and developing new processes for this activity, HECC staff seamlessly coordinated among the many support/engineering entities involved.
- Originally, two shutdowns were scheduled (two transformers per shutdown). By working with the contractor to extend the first shutdown, work was completed in a single event, saving HECC users from the major impact of a second shutdown.

Mission Impact: Scheduled maintenance is a crucial element in sustaining high resource availability for the HECC user community. Careful planning and coordinating helps ensure user impact is kept to a minimum.

Photo of one of four power transformers upgraded in building N233A at NASA's Ames Research Center. The building houses critical high-performance computing resources.

POC: Christopher J. Buchanan, chris.buchanan@nasa.gov, (650) 604-4308, NASA Advanced Supercomputing Division, Computer Sciences Corp.

Increased Building Walk-Through Frequency Pays Off

- HECC Facilities staff doubled the daily building walk-throughs to proactively identify problems.
  - On August 21 (during a walk-through) a puddle of water was discovered in the NAS facility mechanical room basement.
  - Tracing the source of the water, a crack in the bronze-body of a valve was found leaking chilled water under the main computer room into the basement.
- The Facilities team worked closely with ARC engineers to locate the branch shutoff valves, and SGI assessed racks that might lose cooling water during the repair.
- ARC engineers staged repair materials so the repair could be completed quickly without affecting the "cold aisle" temperatures. The racks incurred *no* damage and there was *no* impact to users.
- Quick teamwork between HECC Facilities, SGI, and ARC staff returned the system back to normal by 5:00 pm, eliminating the risk of any damage to Pleiades.

Mission Impact: Improved monitoring of the NASA Advanced Supercomputing facility helps maintain facility up-time and reduces the risk of outages and/or damage to HECC computational resources.

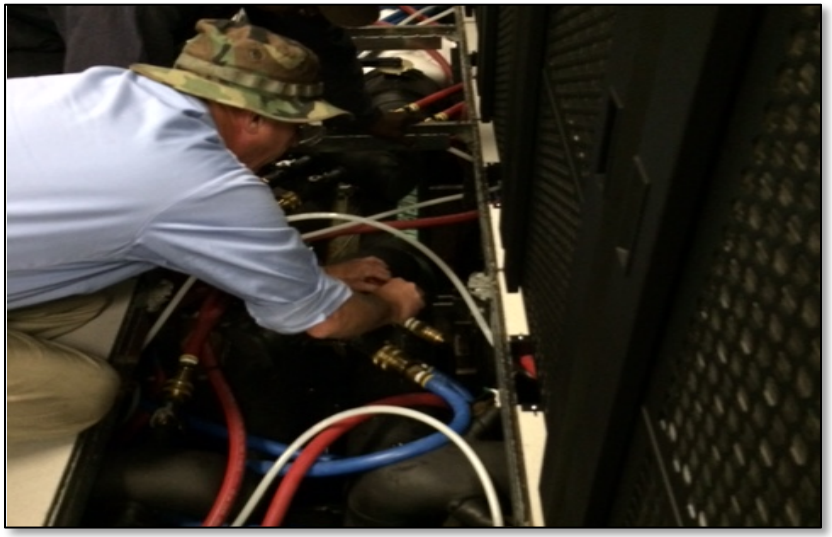

An engineer replaces a chilled-water valve that leaked water into the basement of the NASA Advanced Supercomputing facility.

POCs: Christopher J. Buchanan, chris.buchanan@nasa.gov, 650-604-4308, and Christopher Tanner, christopher.tanner@nasa.gov, (650)604-6754, NASA Advanced Supercomputing Division, Computer Sciences Corp.

Merope Upgraded with Westmere Nodes

- HECC engineers completed the upgrade of 14 racks of Nehalem nodes to Westmere nodes with a live integration of the equipment into the Merope supercomputer.
- Engineers replaced 448 Nehalem nodes with an equivalent number of retired Pleiades Westmere nodes. The upgrade provides an additional 21 teraFLOPS of computing power to Merope.
- During the multi-week process of upgrading Merope, when the hardware was replaced and connected into the InfiniBand fabric, the system remained available to users to run their production workloads.
- By utilizing repurposed Pleiades nodes, HECC can provide additional SBUs to users at minimal cost to the agency.

Mission Impact: Repurposing retired hardware enables HECC to deliver additional computational cycles to NASA users.

The Merope supercomputer now consists of 1,152 Westmere nodes that were retired from Pleiades. These nodes continue to provide computational cycles to HECC users in the secondary compute facility at NASA Ames.

POCs: Bob Ciotti, bob.ciotti@nasa.gov, (650) 604-4408, NASA Advanced Supercomputing Division; Davin Chan, davin.chan@nasa.gov, (650) 604-3613, NASA Advanced Supercomputing Division, Computer Sciences Corp.

Vendor Selected to Augment Lustre Storage Capacity

- HECC selected an additional Lustre storage solution to augment the existing Lustre filesystems, which nearly doubles the short-term storage capacity available on HECC resources.
- The selected solution was submitted by SGI/ NetApp, and provided the best value for the performance and capacity.
- The solution will provide 18.7 petabytes (PB) of storage and will be capable of over 80 gigabytes/second (GB/s) of random I/O performance and over 100 GB/s of sequential performance.
- The storage delivery is planned for late September, and will be released for production in October after completing acceptance testing and validation.

Mission Impact: Increased storage capacity and performance will enable HECC users to more fully utilize the computing resources of Pleiades and Endeavour.

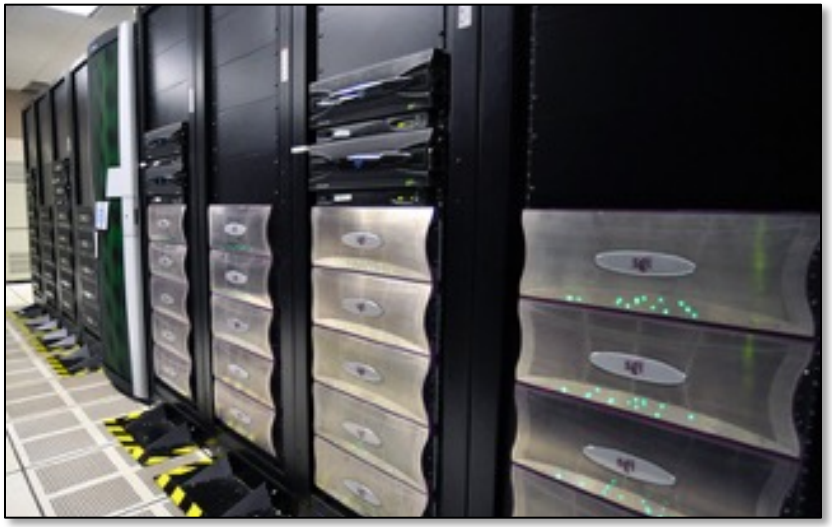

The current Lustre storage system consists of DataDirect Networks (DNN) and NetApp storage that totals approximately 20 petabytes of storage. The new storage will nearly double the capacity, with an additional 18.7 PB.

POCs: Bob Ciotti, bob.ciotti@nasa.gov, (650) 604-4408, NASA Advanced Supercomputing Division; Davin Chan, davin.chan@nasa.gov, (650) 604-3613, NASA Advanced Supercomputing Division, Computer Sciences Corp.

HECC Provides Dedicated Support to ERA Airframe Noise Reduction Simulations

- HECC provided dedicated support to NASA's Environmentally Responsible Aviation (ERA) project, "Evaluation of Noise Reduction Concepts for Flaps," prior to the June 1–9, 2015 annual facility maintenance shutdown, in order to help the project meet critical milestones due at the end of June.
- HECC provided the project exclusive access to 16 Ivy Bridge racks (23,040 cores total) from April 5 to May 31 for performing multiple, large-scale coarse- and medium-resolution noise-reduction simulations for a full-scale aircraft. Without these dedicated resources, the project would not have been able to complete most of these simulations needed for the milestones.
- HECC teams, together with SGI support staff, provided concierge service to address hardware and software issues either created or encountered by ERA project team. Such efforts resulted in improved batch job monitoring and higher stability in the Pleiades supercomputer's network and file systems.
- The project team plans to repeat the study using finer-resolution grids and extending the simulations to several flap and landing gear configurations.

Mission Impact: Simulations performed on the Pleiades supercomputer represent a critical element of NASA's Environmentally Responsible Aviation project, in the evaluation and maturation of effective noise-reduction concepts for aircraft flaps and landing gear.

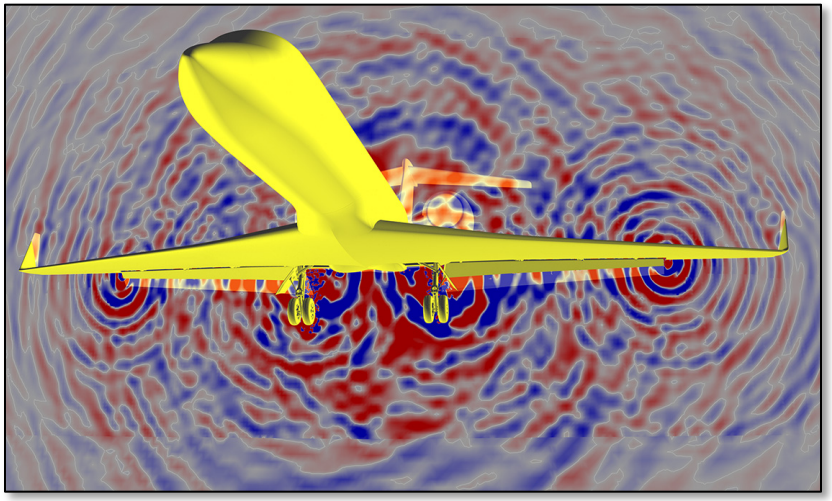

Simulated radiated sound field produced by a full-scale Gulfstream aircraft during landing, with flaps and main landing gear deployed. *Ehab Fares, Exa Corporation; Patrick Moran, NASA/Ames*

POCs: Sherry Chang, sherry.chang@nasa.gov, (650) 604-1272, Computer Sciences Corp.; Catherine Schulbach, catherine.h.schulbach@nasa.gov, (650) 604-3180; NASA Advanced Supercomputing Division

HECC Staff Make Major Contributions to Improving SGI MPT Library

- The Application Performance and Productivity (APP) group helped SGI identify defects in several released and pre-release alpha versions of their Message Passing Toolkit (MPT) package that supports parallel programming across HECC's networked systems.
- With Pleiades' extensive and constantly evolving InfiniBand network and Lustre I/O fabrics, HECC and SGI staff regularly address communication issues with new methods implemented in the MPT library for improving Pleiades' stability.
  - Defects introduced in recent MPT versions caused performance regressions, program aborts, and program "hangs" in various user applications.
  - APP consultants researched the characteristics, created simple reproducers, and found appropriate test cases for each user-reported problem, and worked with SGI engineers through several iterations of patches to fix all known problems.
- Weekly tracking of job failures now shows a significant decrease in jobs failures due to MPT after users migrated to the new and improved MPT 2.12r26 library on August 3.

Mission Impact: Improving the performance and reliability of the SGI MPT library has a huge benefit to productivity for users of the Pleiades supercomputer, resulting in fewer job failures.

Improvements to the SGI Message Passing Toolkit software allows for faster communication by reducing congestion on the InfiniBand network on Pleiades.

POCs: Sherry Chang, sherry.chang@nasa.gov, (650) 604-1272; Johnny Chang, johnny.chang@nasa.gov, (650) 604-4356, NASA Advanced Supercomputing Division, Computer Sciences Corp.

Shift Transfer Tool Enhanced With Automatic Transport TCP Tuning

- Shift is an in-house automated data transfer tool that automatically encapsulates HPC best practices for increased usability, reliability, and performance.
- HECC Systems staff measured the performance of several remote transports across high-speed wide area network (WAN) and local area network (LAN) links, while varying Transmission Control Protocol (TCP) window size and the number of TCP streams.
- Results showed that the fixed values previously used within Shift for bbftp (4-megabyte window size and 4 streams) were non-optimal, in some cases.
  - WAN performance improved over 3x from 236 megabytes per second (MB/s) to 755 MB/s by increasing the TCP window size to 64 MB on high-speed links.
  - LAN performance improved almost 2x from 317 MB/s to 624 MB/s by reducing the number of parallel TCP streams down to a single stream.
- Systems staff augmented Shift with the ability to adjust the number of streams based on transfer type (WAN/ LAN) and use a larger window size when a high-speed WAN link is detected.
  - This capability has been implemented and will be deployed during the next Shift release cycle.
  - Users will transparently achieve higher transfer rates.

Mission Impact: By automatically adjusting Transmission Control Protocol settings based on the specific characteristics of each data transfer, HECC's Shift tool maximizes transfer performance and reduces user turnaround time.

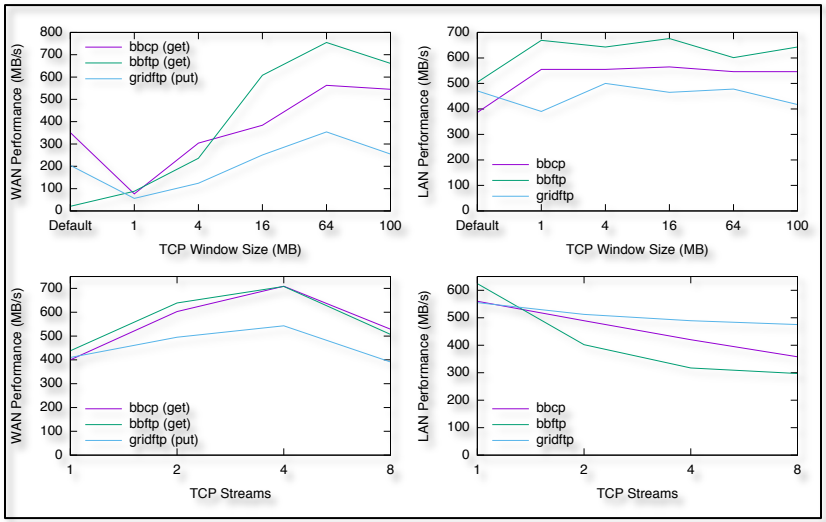

Different transport settings for the TCP window size (top) and for the number of parallel TCP streams (bottom) can significantly impact data transfer performance, with the optimum values for each differing between high speed WAN (left) and LAN (right) links.

POCs: Bob Ciotti, bob.ciotti@nasa.gov, (650) 604-4408, NASA Advanced Supercomputing Division; Davin Chan, davin.chan@nasa.gov, (650) 604-3613, NASA Advanced Supercomputing Division, Computer Sciences Corp.

Space Weather Simulations for Improving Solar Storm Prediction

- With the goal of advancing space weather forecasting, researchers at San Diego Supercomputing Center (SDSC) performed the most detailed magnetosphere/ solar wind interactions simulations ever run.
- Enabled by the Pleiades supercomputer, the SDSC team solved several science mysteries:
  - Identified conditions where parcels of solar wind plasma are able to reach deep inside the Earth's magnetosphere.
  - These parcels of plasma impinging on the magnetopause (the boundary of the Earth's magnetosphere) cause space weather effects.
  - identified (and demonstrated for the first time) the mechanism for the formation of large-scale vortices inside the magnetosheath.
- Recent comparisons of SDSC simulation data matched well with direct observations from the European Space Agency's Cluster spacecraft. Researchers continue to compare the SDSC predictions with observations from both NASA and ESA spacecraft.
- Petascale, kinetic simulations on massively parallel systems such as Pleiades enabled SDSC to exceed previous state-of-the-art magnetospheric models.

Mission Impact: Enabled by the Pleiades supercomputer, these high-fidelity simulations will improve the accuracy of space weather modeling and forecasting to help mitigate costly solar storm damage.

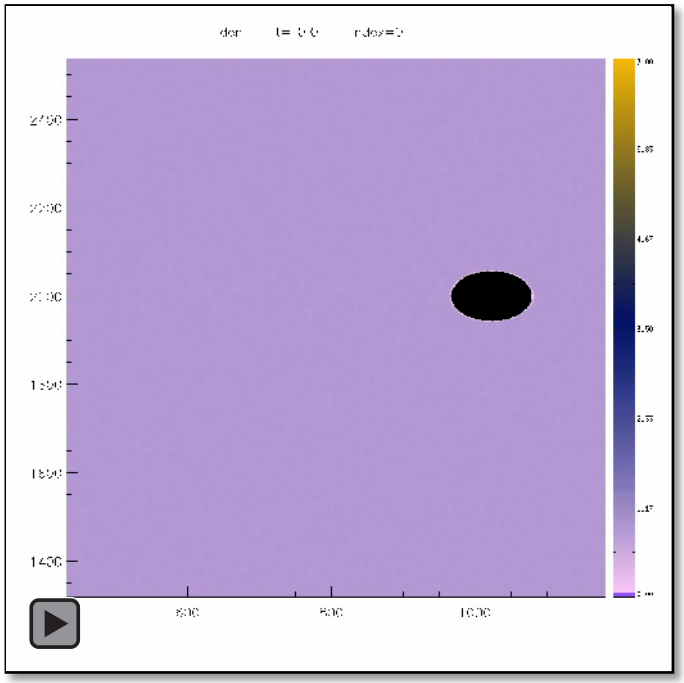

A global hybrid simulation of Earth's magnetosphere, showing for the first time the formation of high-velocity jets of plasma penetrating and reaching the Earth's magnetopause. The jets cause space weather effects.

POC: Homa Karimabadi, homakar@gmail.com, University of California, San Diego/SciberQuest, Inc.

** HECC provided supercomputing resources and services in support of this work*

Overset Grid Simulations of the Open Rotor Configuration Using LAVA

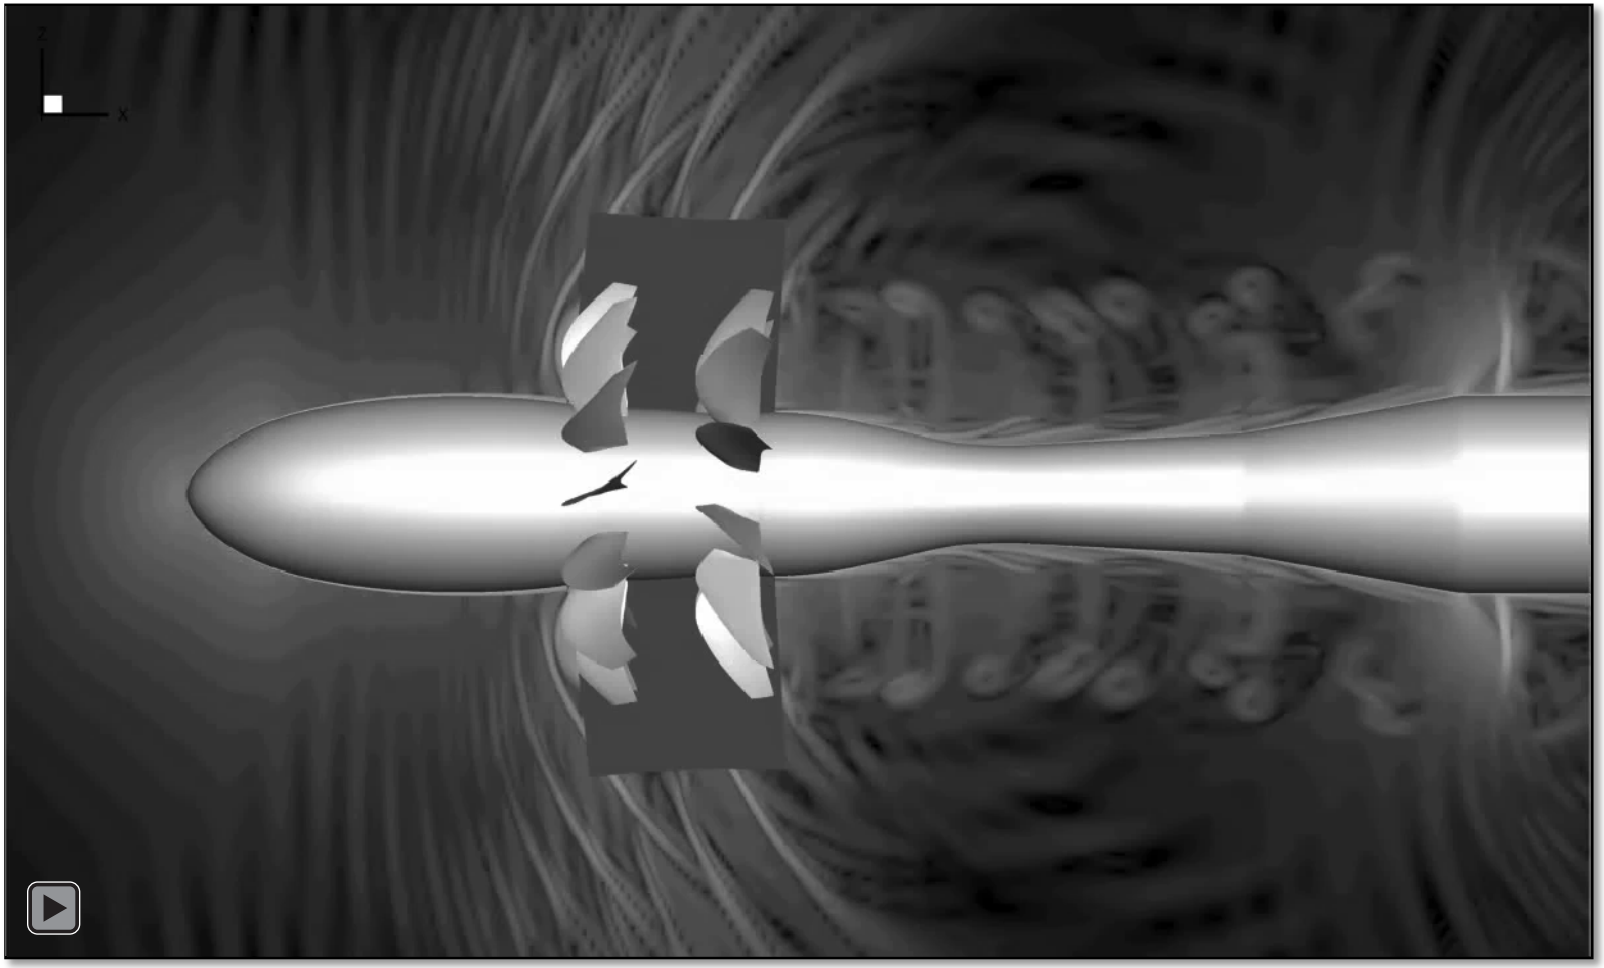

Contour plot of the magnitude of density gradient highlighting the acoustic wave generation from the blades, the wake from the blades, and the interaction of the wake with the fish tail shock located on the downstream side of the hub.
***POC**: Jeffrey Housman, jeffrey.a.housman@nasa.gov, (650) 604-5455, NASA Advanced Supercomputing Division.*

HECC Facility Hosts Several Visitors and Tours in August 2015

- HECC hosted 15 tour groups in August; guests learned about the agency-wide missions being supported by HECC assets, some of the groups also viewed the D-Wave 2x quantum computer system. Visitors this month included:
  - Participants from NASA Aeronautics Research Institute's Transformative Vertical Flight Workshop received a presentation from research scientist Neal Chaderjian, who showed his advanced rotorcraft wake computational fluid dynamics visualizations.
  - Sally-Ann Watts, Australian Consul General.
  - A group of 20 NASA Space Launch System aerospace specialists, who were attending an Ames Technical Information Meeting.
  - Congressman Don Beyer from Virginia.
  - Jean-Francois Barthelemy, Langley Research Center Chief Technologist.
  - A large group of attendees from the NASA Cost Symposium, including Lisa Ziehmann, NASA Deputy Chief Financial Officer, Finance and Systems; Joe McIntyre, NASA Associate Deputy Chief Financial Officer, Finance; Adam Bethon, Financial Management Analyst, and 50 others.
  - Mark Mozena, Congressman Mike Honda's key staff representative on space/science issues.
  - A large group from Altair, who attended a local PBS (Portable Batch System) User Group meeting.

Nick Bonifas (pink shirt), of the Publications and Media group at the NASA Advanced Supercomputing facility, led part of the NASA Cost Symposium group through the main computer room housing the Pleiades supercomputer.

POC: Gina Morello, gina.f.morello@nasa.gov, (650) 604-4462, NASA Advanced Supercomputing Division

Papers

- "**G2 and Sgr A*: A Cosmic Fizzle at the Galactic Center**," B. Morsony, B. Gracey, J. Workman, D.-S. Yoon, arXiv:1508.00384 [astro-ph.HE], August 3, 2015. * *http://arxiv.org/abs/1508.00384*
- "**Solar Wind Interaction with the Martian Upper Atmosphere: Crustal Field Orientation, Solar Cycle and Seasonal Variations**," C. Dong, et al., Journal of Geophysical Research: Space Physics, August 14, 2015. * *http://onlinelibrary.wiley.com/doi/10.1002/2015JA020990/full*
- "**A Simulation Study of the Relationship Between Tail Dynamics and the Aurora**," M. Ashour-Abdalla, chapter in *Magnetospheric Plasma Physics: The Impact of Jim Dungey's Research*, volume 41 of the Astrophysics and Space Sciences Proceedings series, August 2015. * *http://link.springer.com/chapter/10.1007/978-3-319-18359-6_5*

** HECC provided supercomputing resources and services in support of this work*

News and Events

- **Behind the Scenes at NASA's Computer Powerhouse**, *Mountain View Voice*, August 28, 2015—Piyush Mehrotra, chief of the NASA Advanced Supercomputing Division at NASA Ames Research Center, discusses how the agency uses the Pleiades supercomputer for Earth and space science, aeronautics research, and next-gen space launch vehicles.
  http://mv-voice.com/news/2015/08/27/behind-the-scenes-at-nasas-computer-powerhouse

HECC Utilization

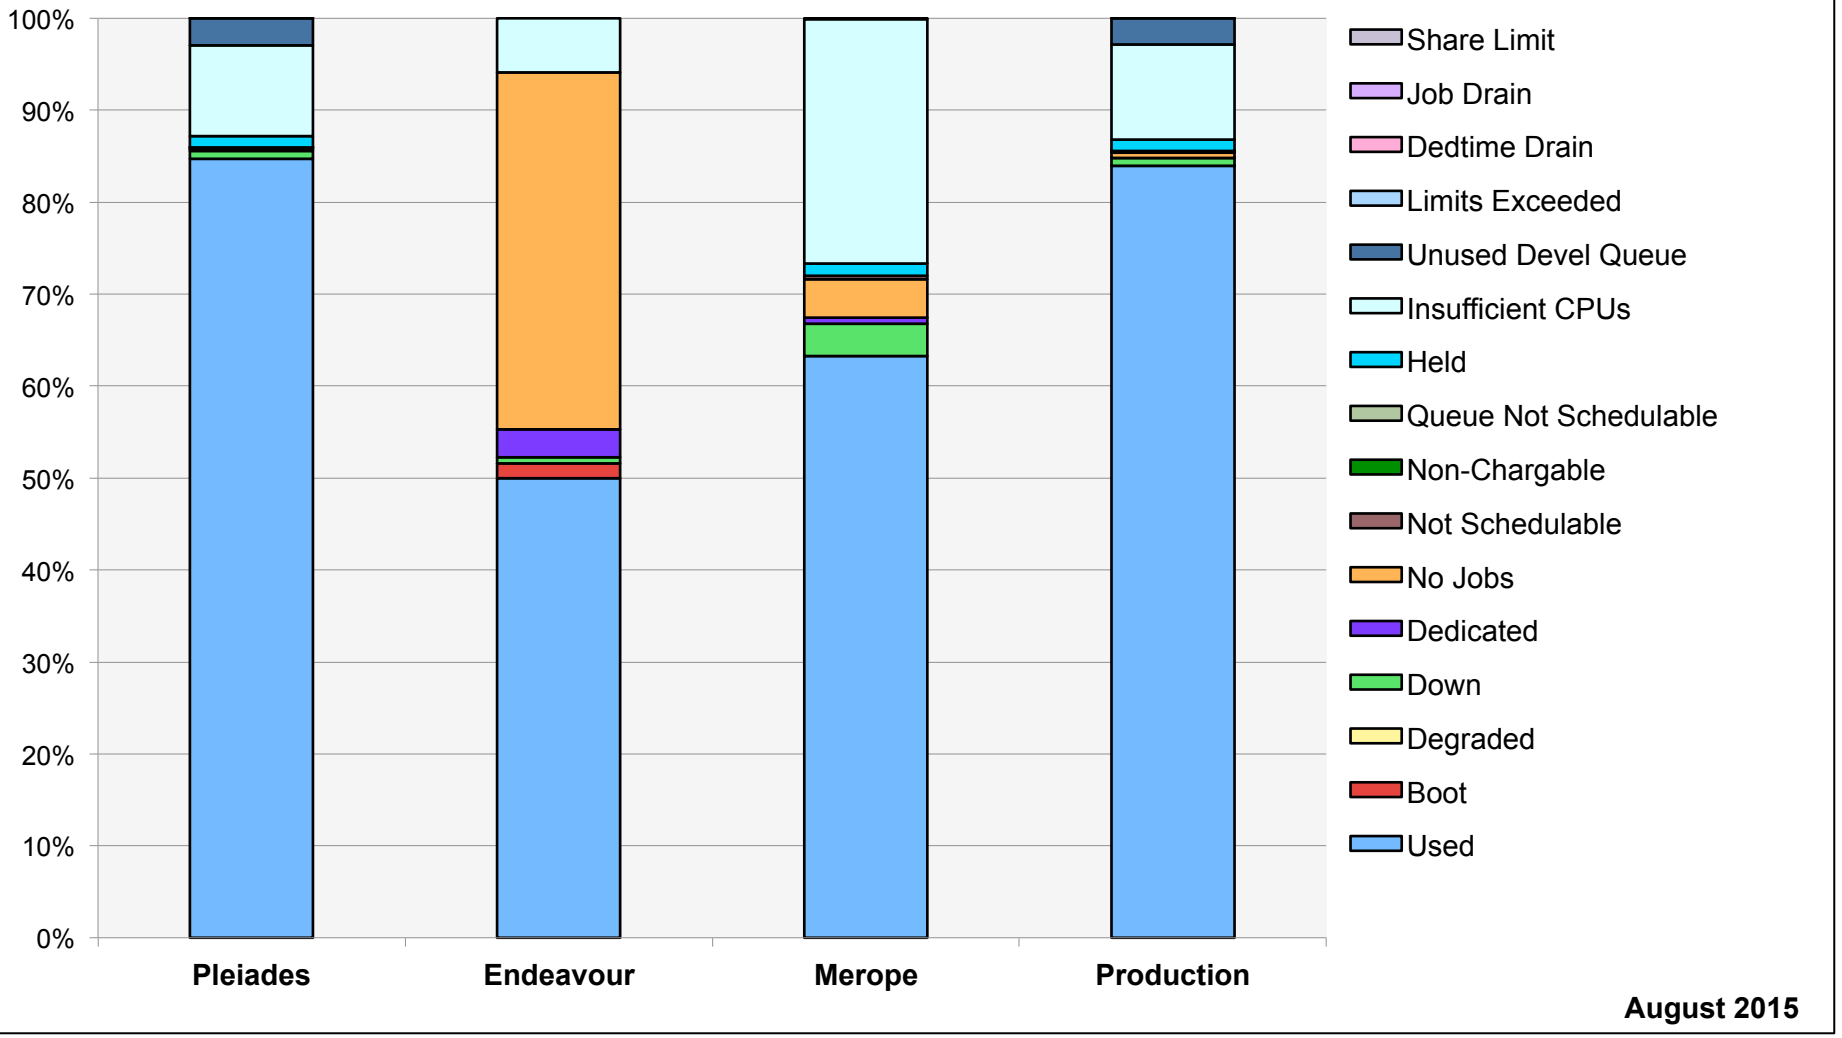

HECC Utilization Normalized to 30-Day Month

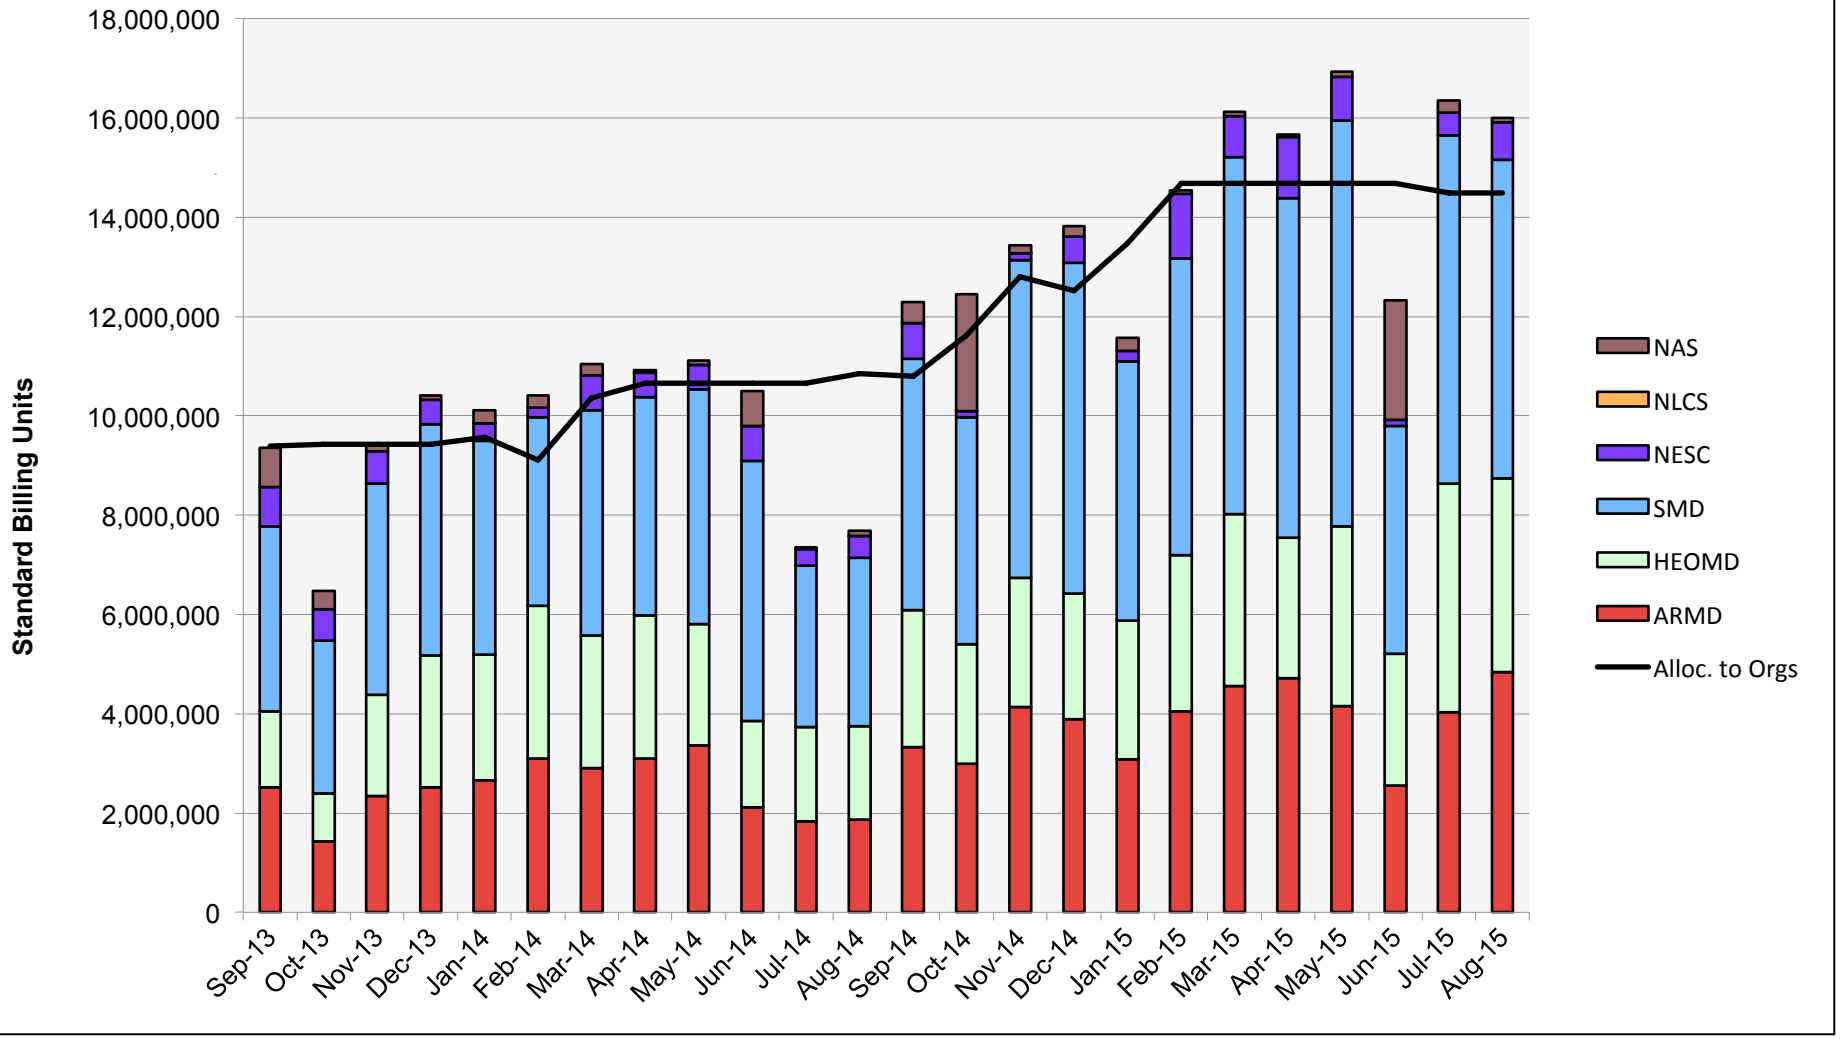

HECC Utilization Normalized to 30-Day Month

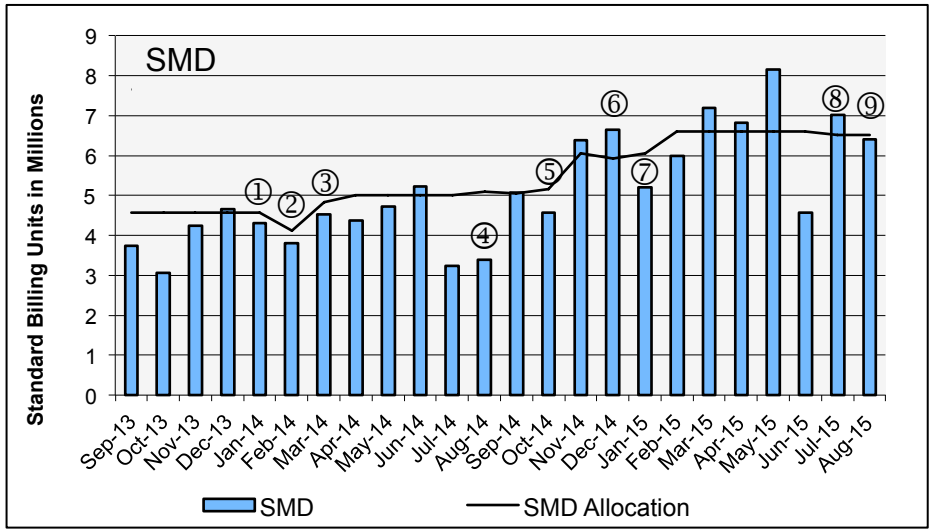

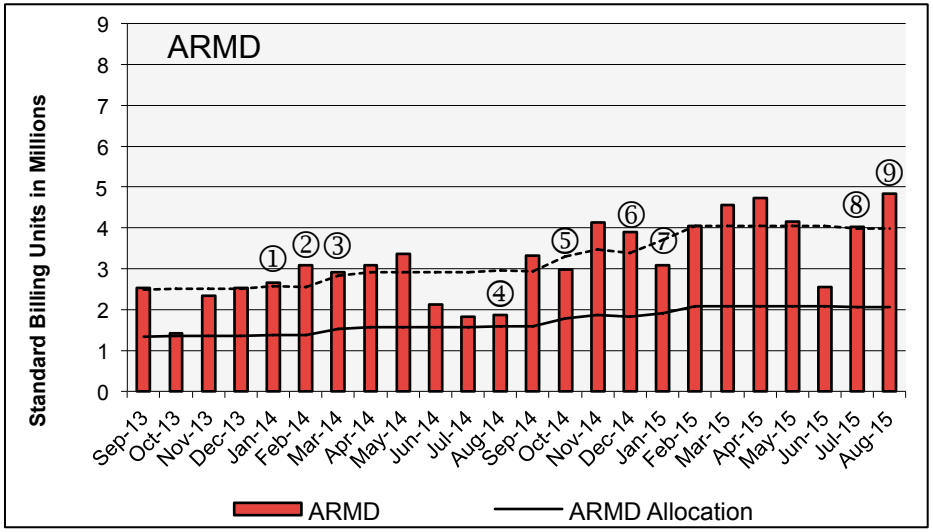

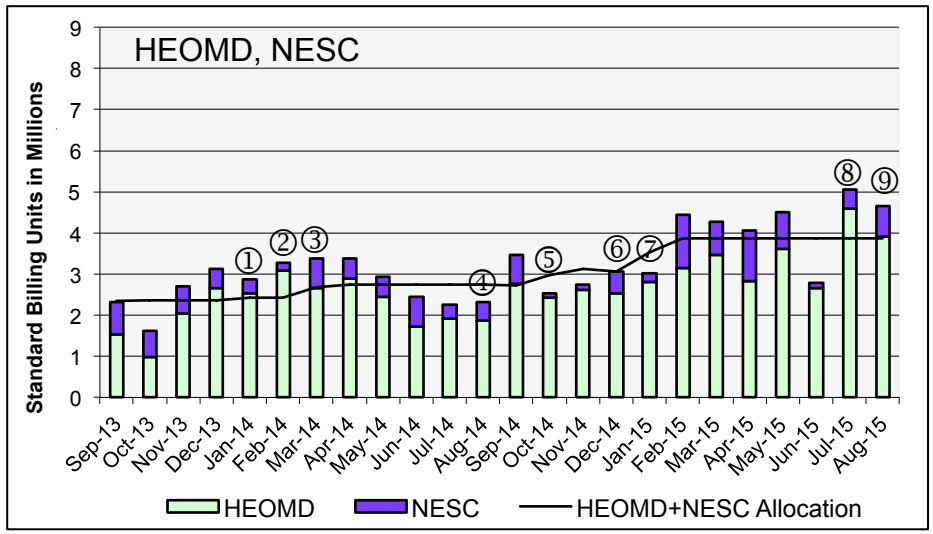

① 6 Ivy Bridge Racks added; 20 Nehalem, 12 Westmere Racks Retired

② 8 Ivy Bridge Racks added mid-Feb; 8 Ivy Bridge Racks added late Feb.

③ 4 Ivy Bridge Racks added mid-March

④ 6 Westmere Racks added to Merope, Merope Harpertown retired

⑤ 16 Westmere Racks retired; 10 Nehalem Racks and 2 Westmere Racks added to Merope; 3 Ivy Bridge Racks added; 15 Haswell Racks added

⑥ 16 Westmere Racks retired

⑦ 14 Haswell racks added

⑧ 7 Merope Nehalem Racks removed

⑨ 7 Merope Westmere Racks added

Tape Archive Status

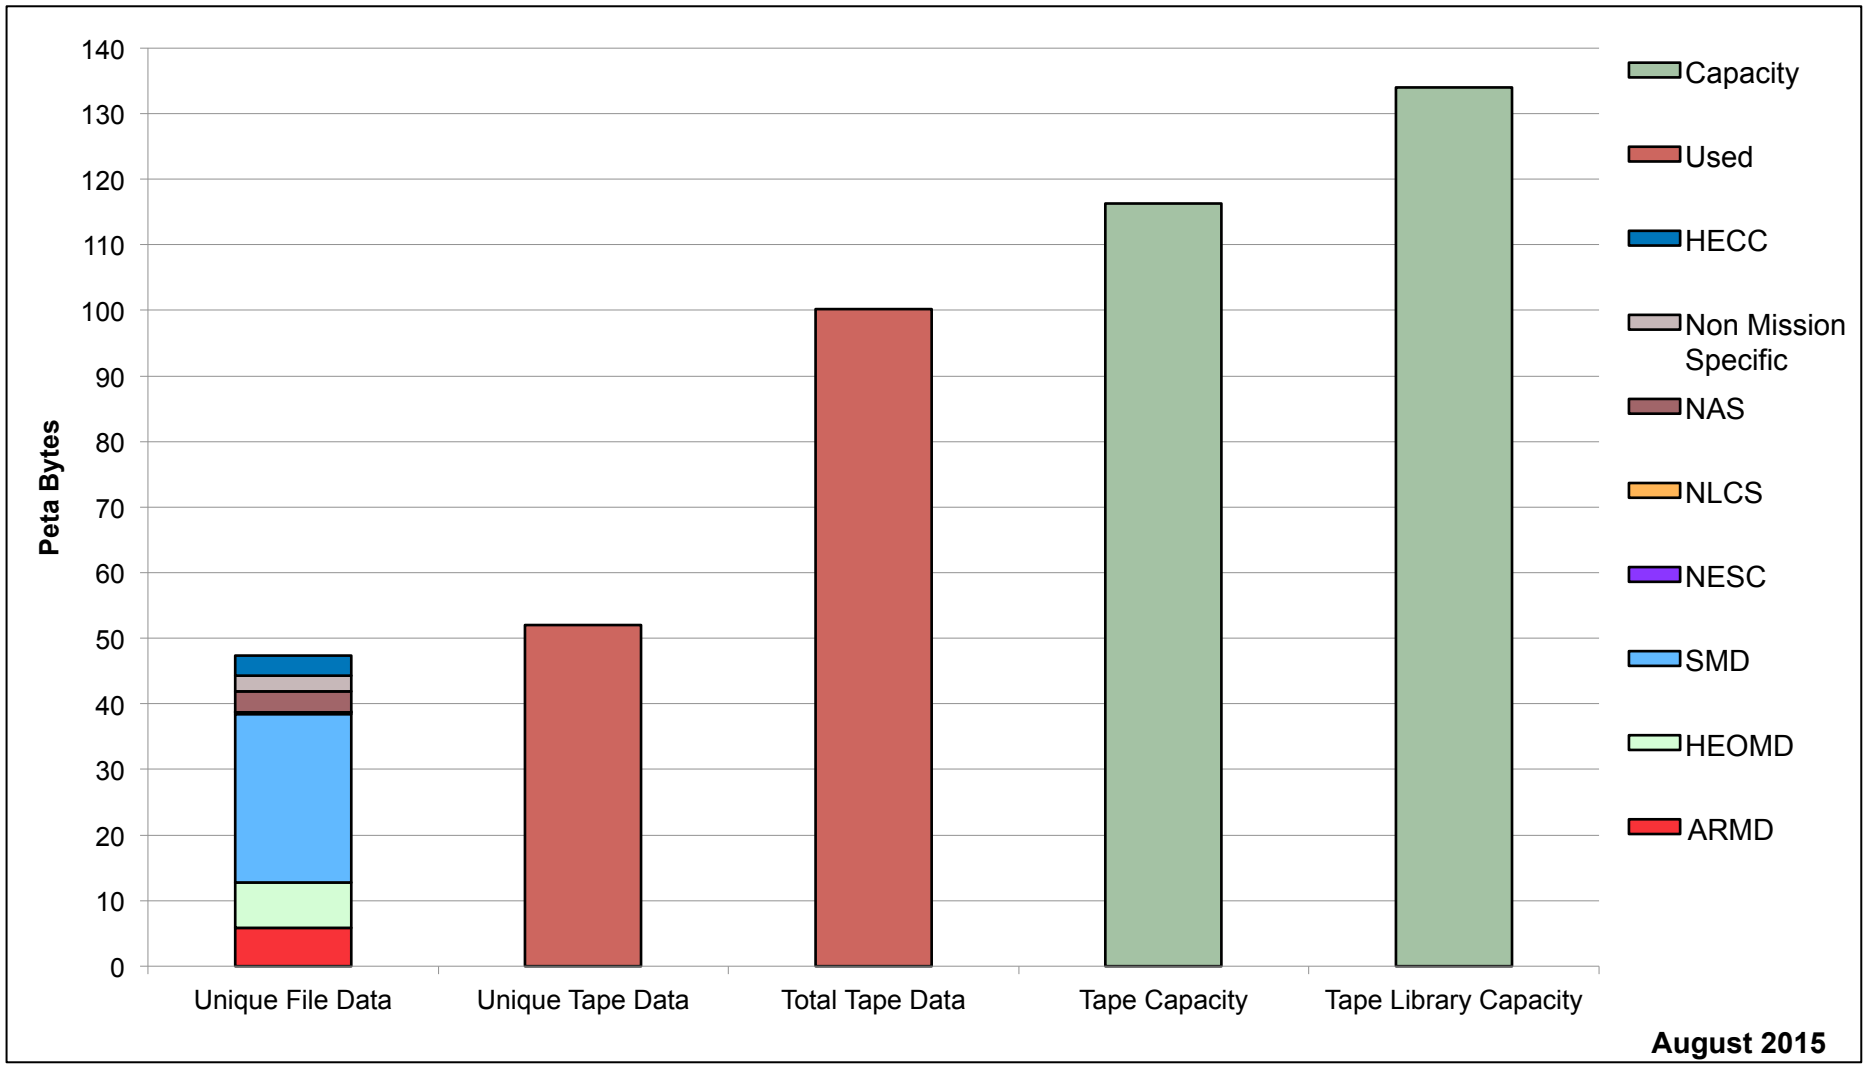

Tape Archive Status

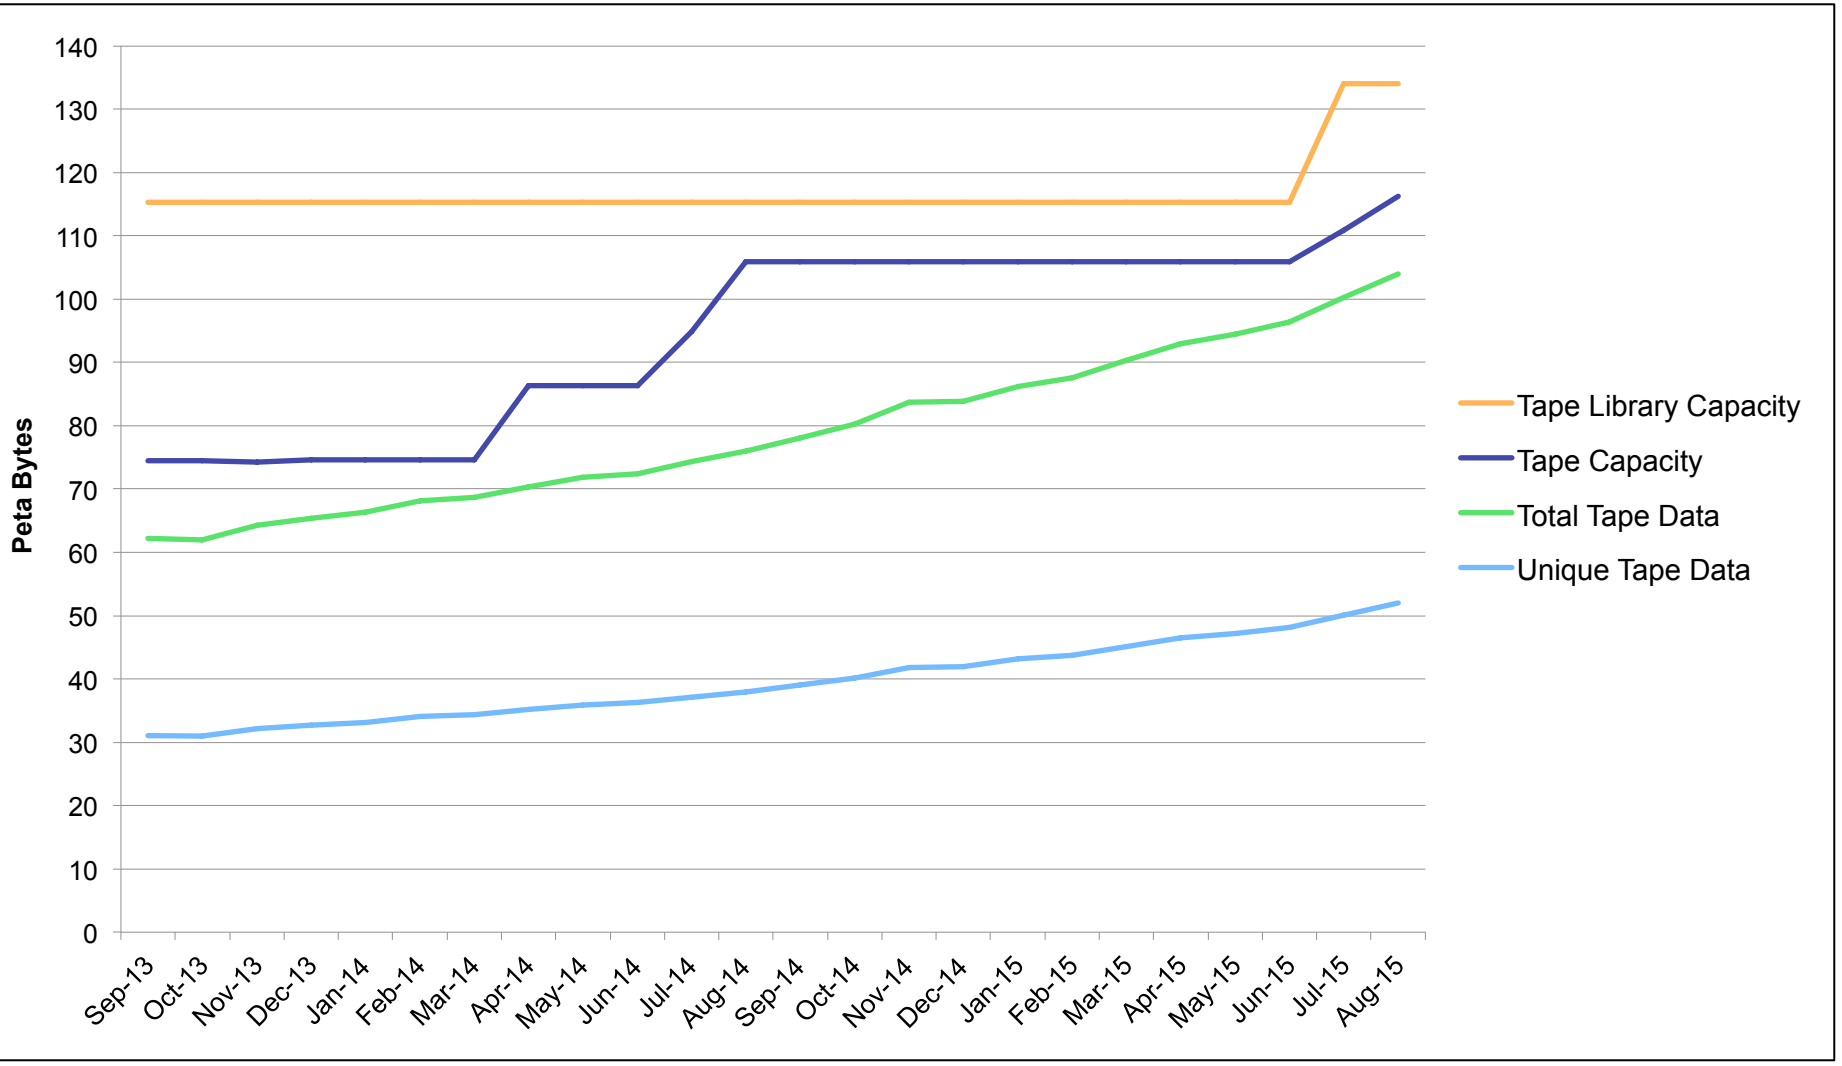

Pleiades:
SBUs Reported, Normalized to 30-Day Month

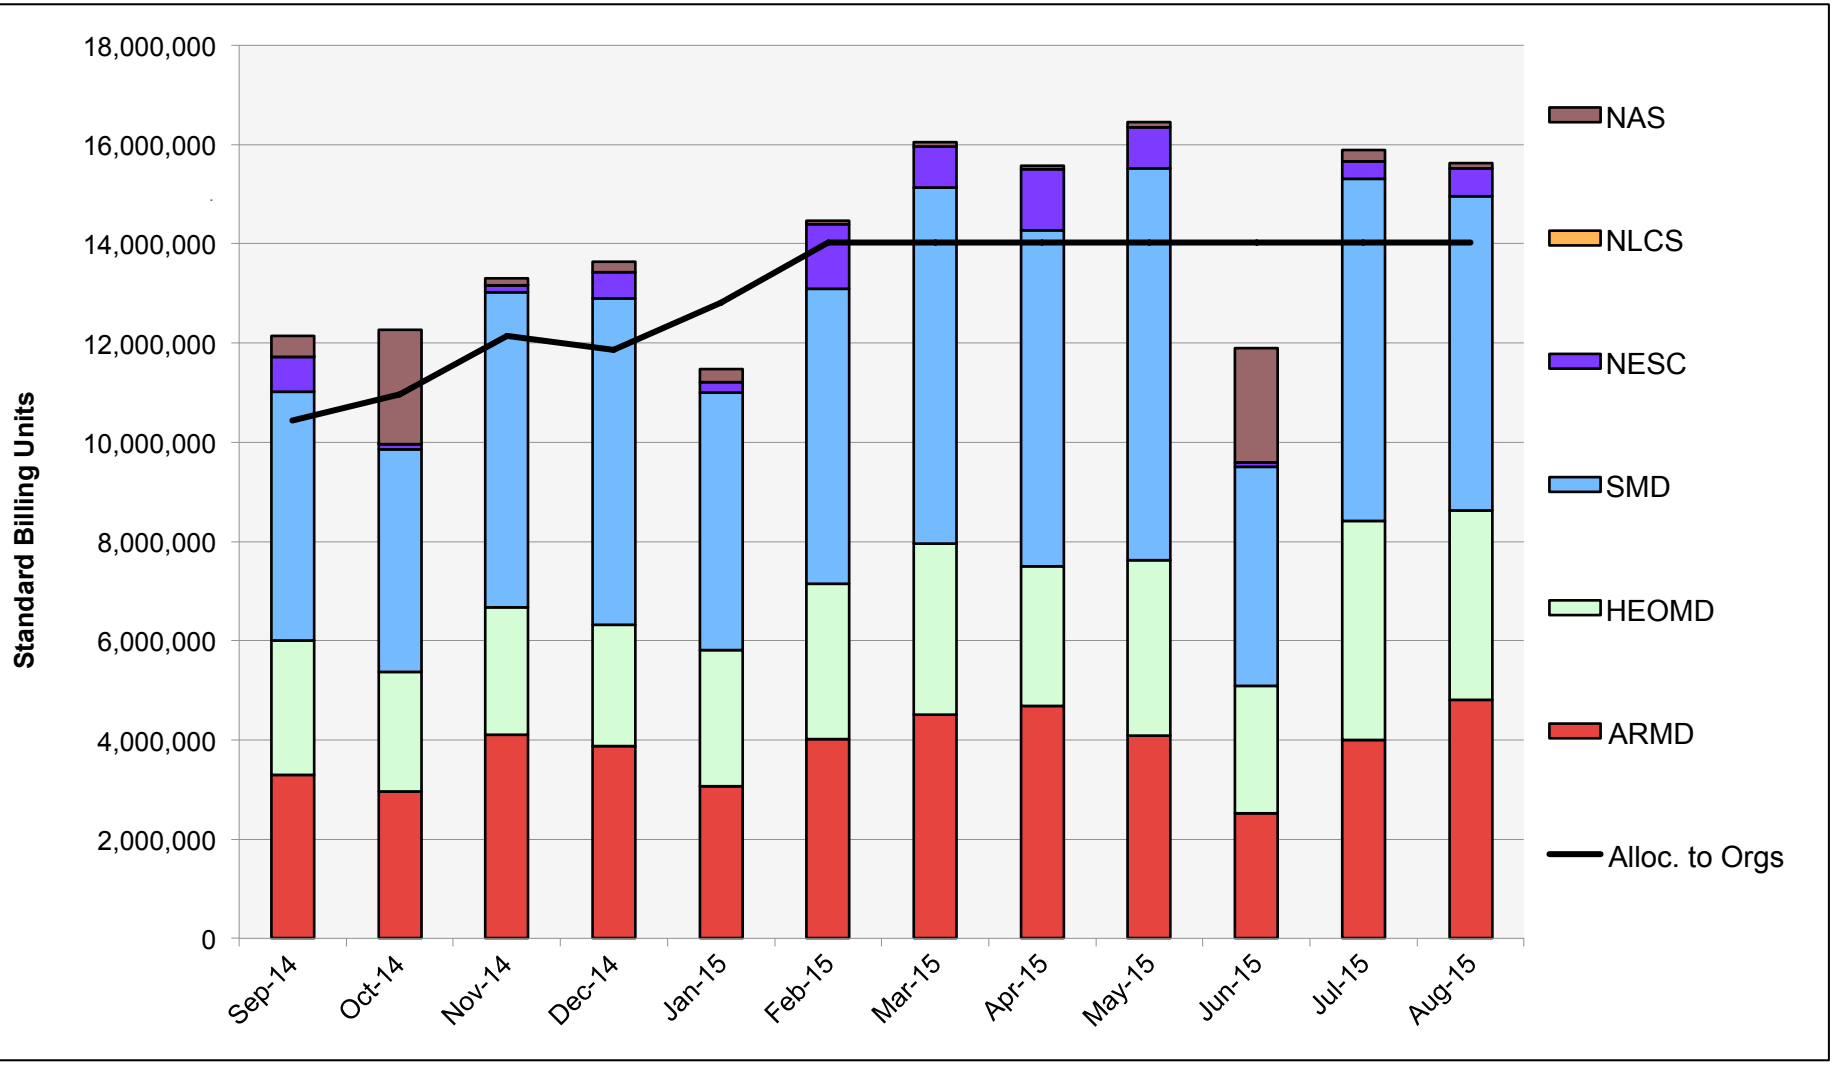

Pleiades:
Devel Queue Utilization

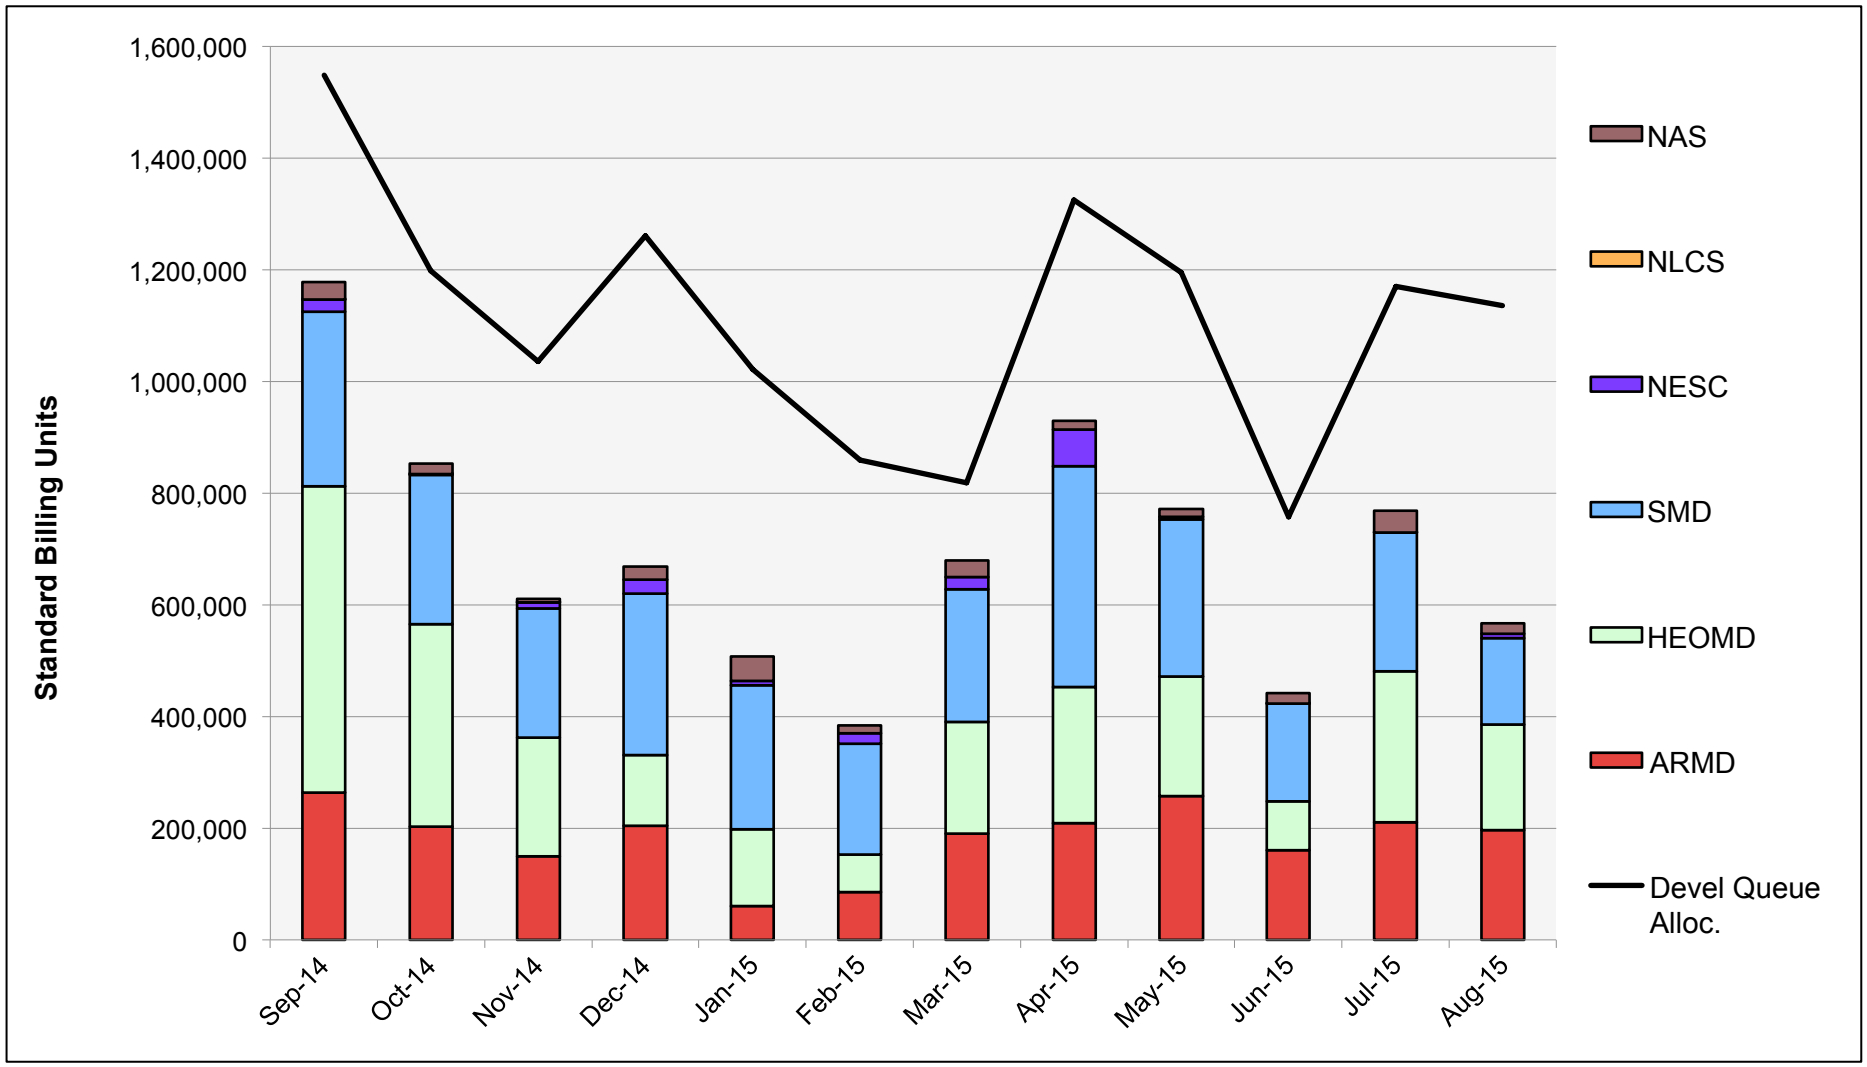

Pleiades:
Monthly Utilization by Job Length

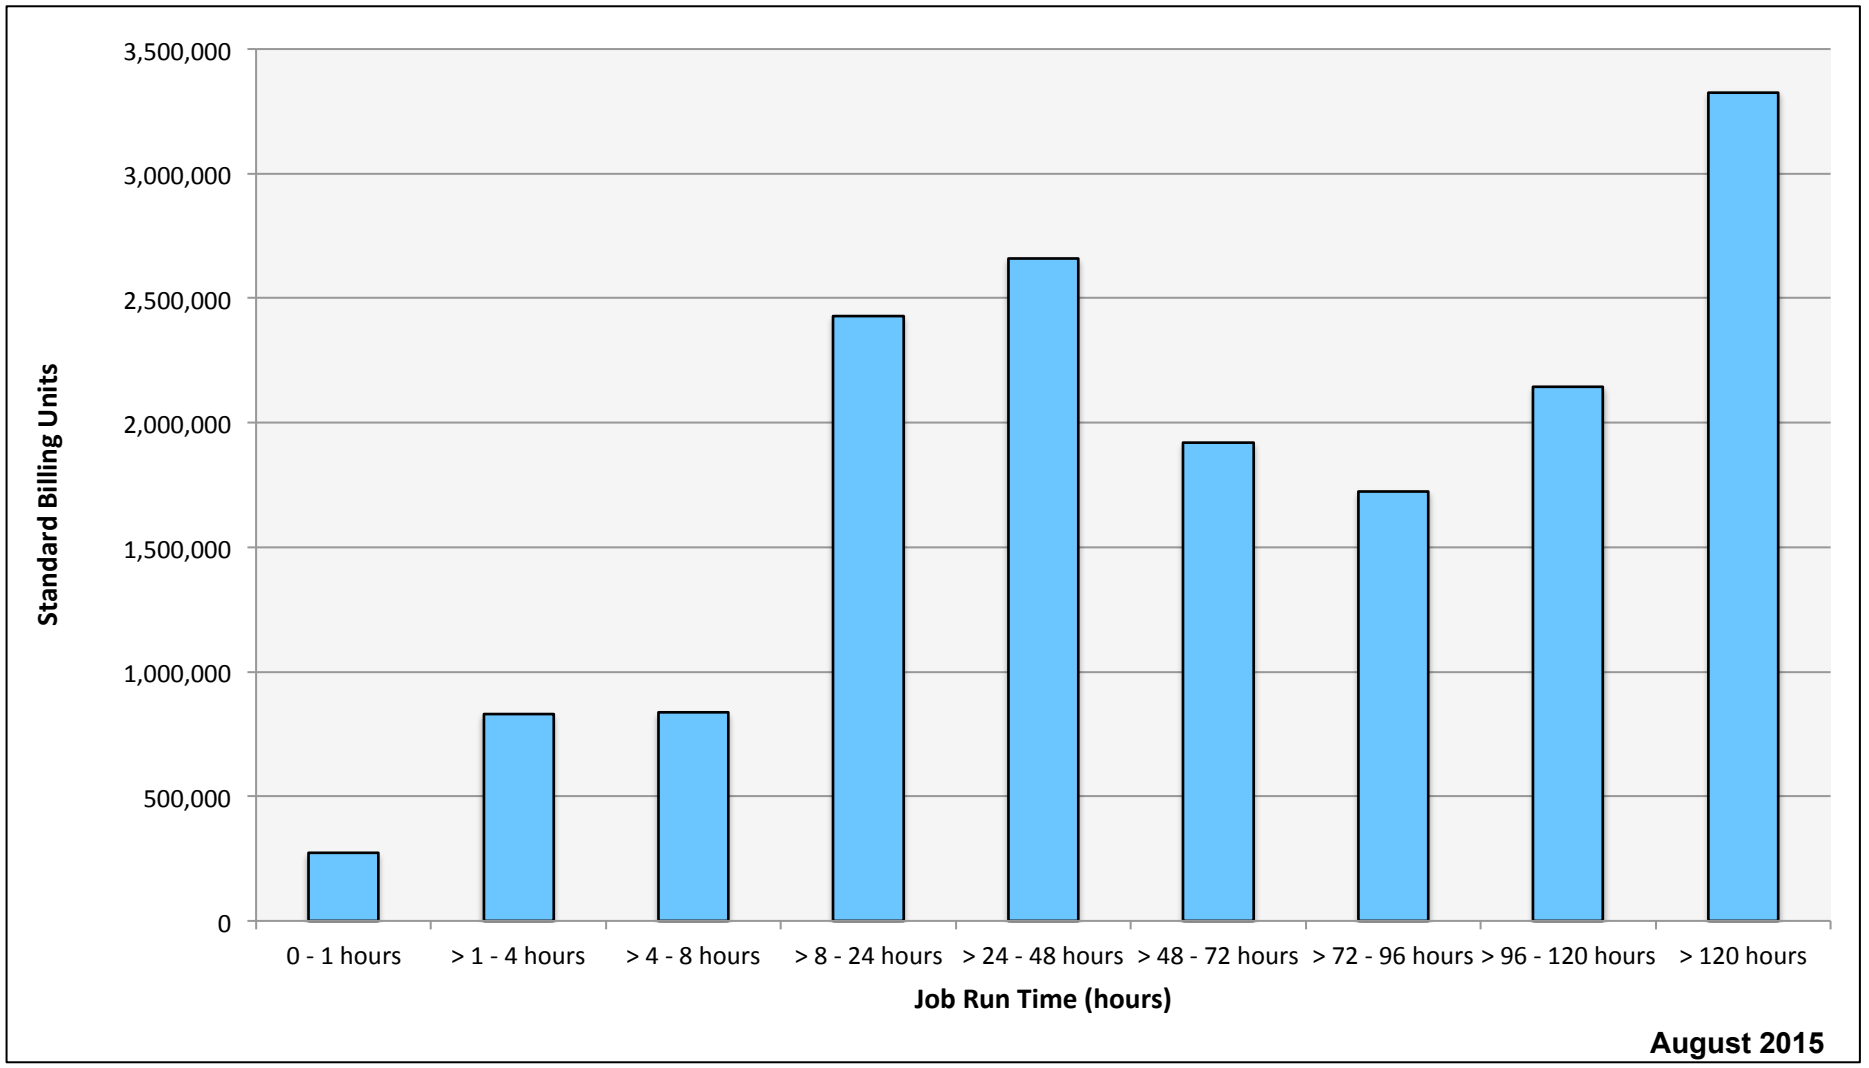

Pleiades:
Monthly Utilization by Size and Mission

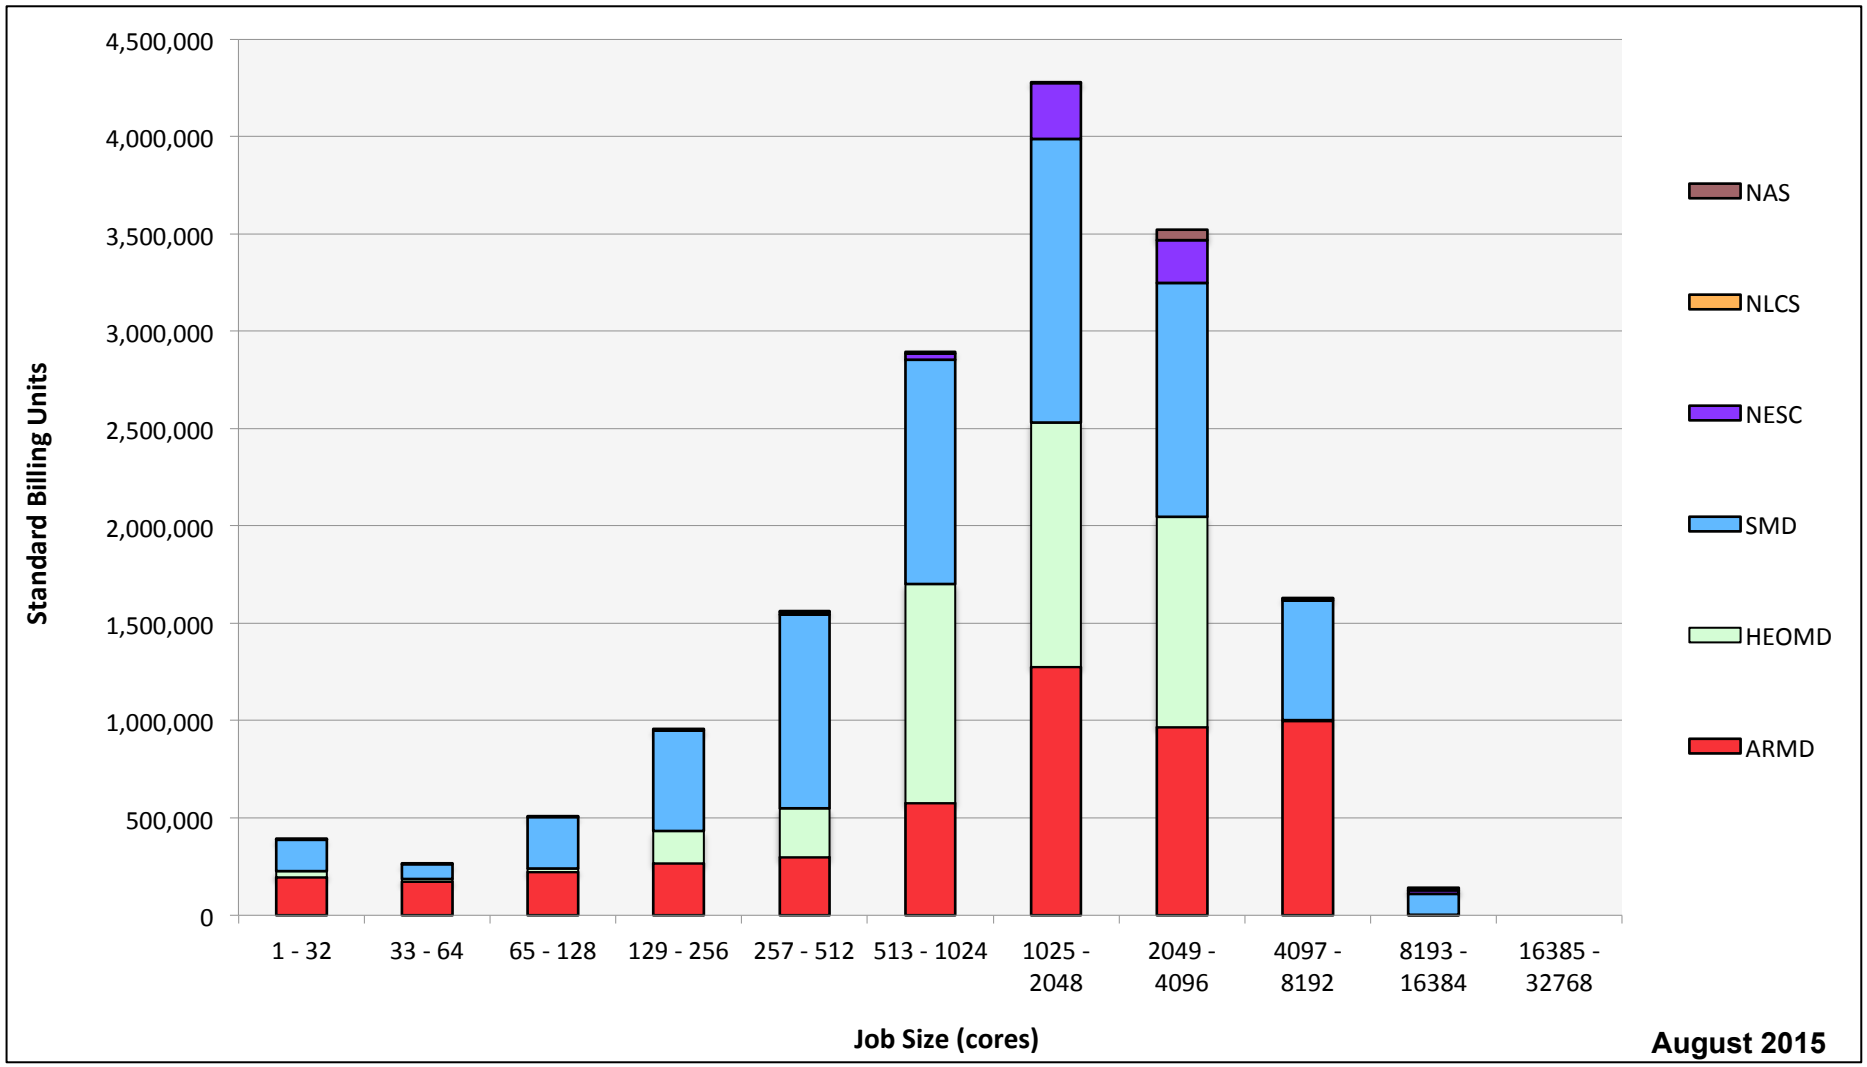

Pleiades:
Monthly Utilization by Size and Length

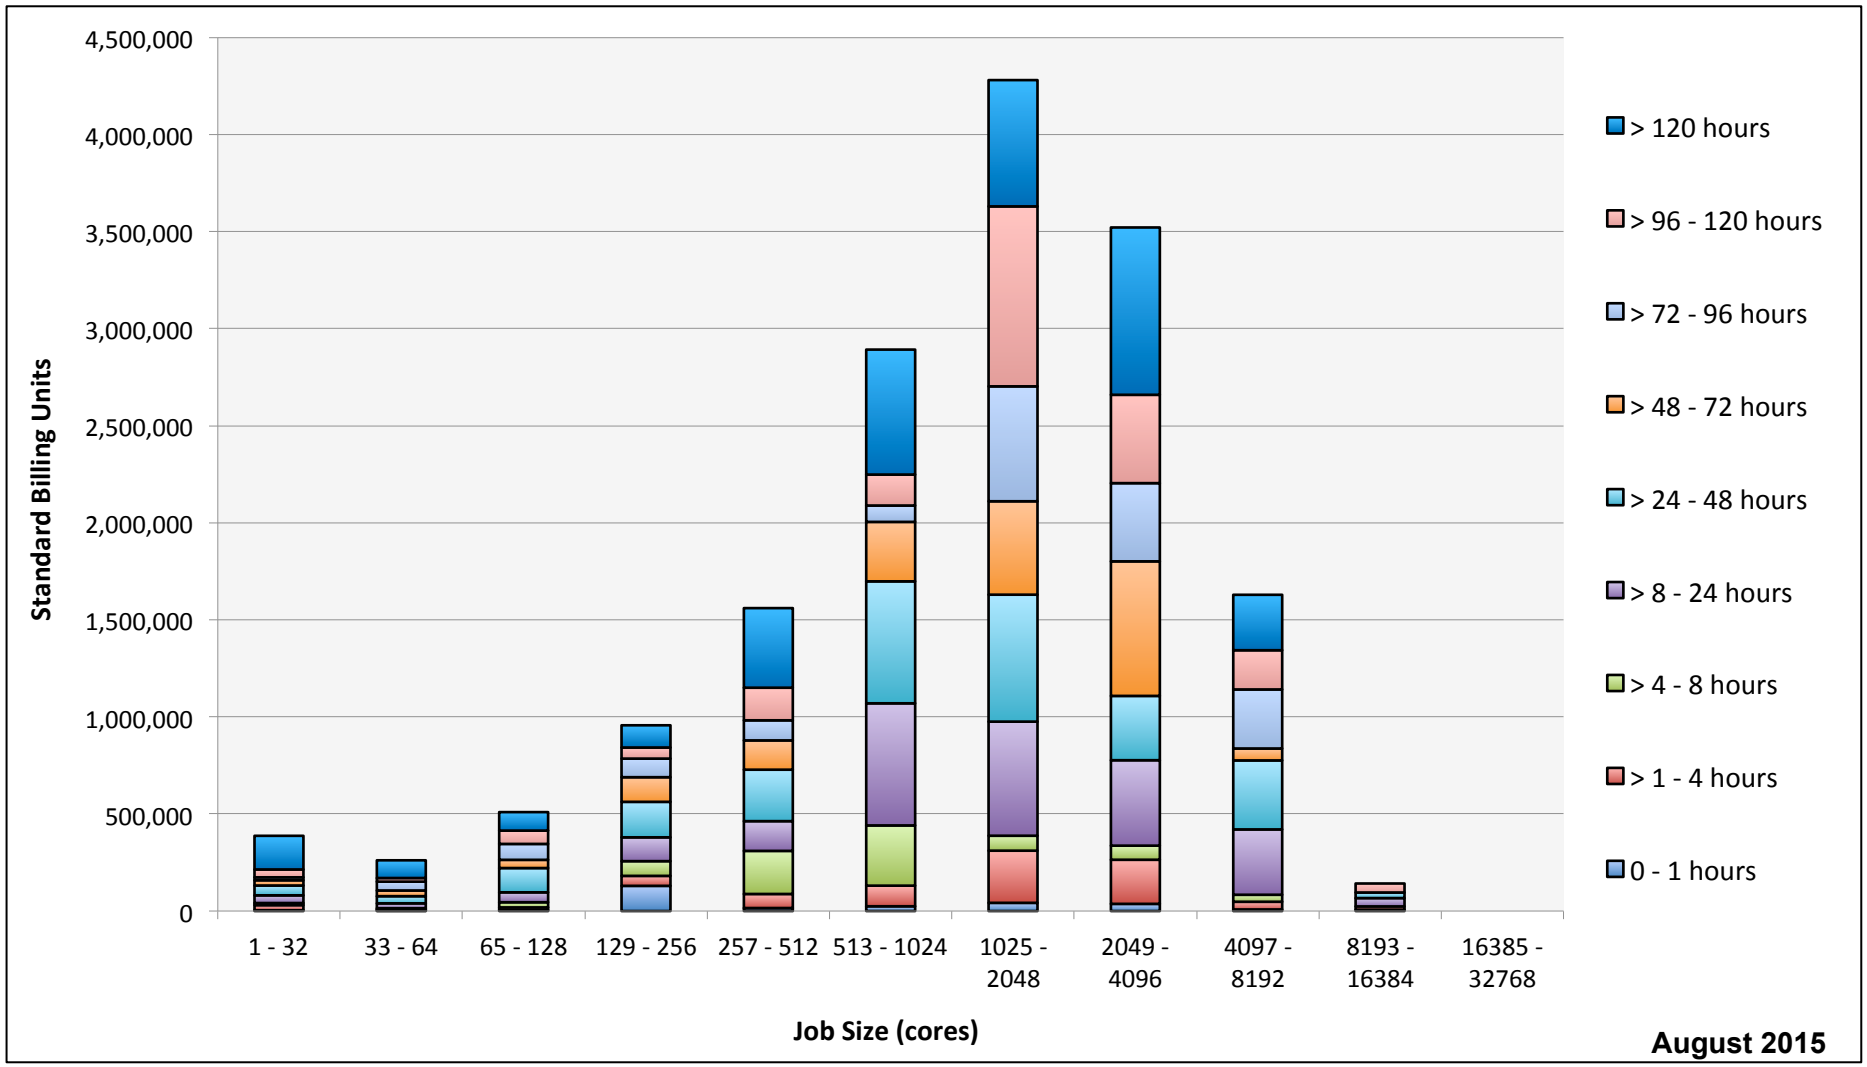

Pleiades:
Average Time to Clear All Jobs

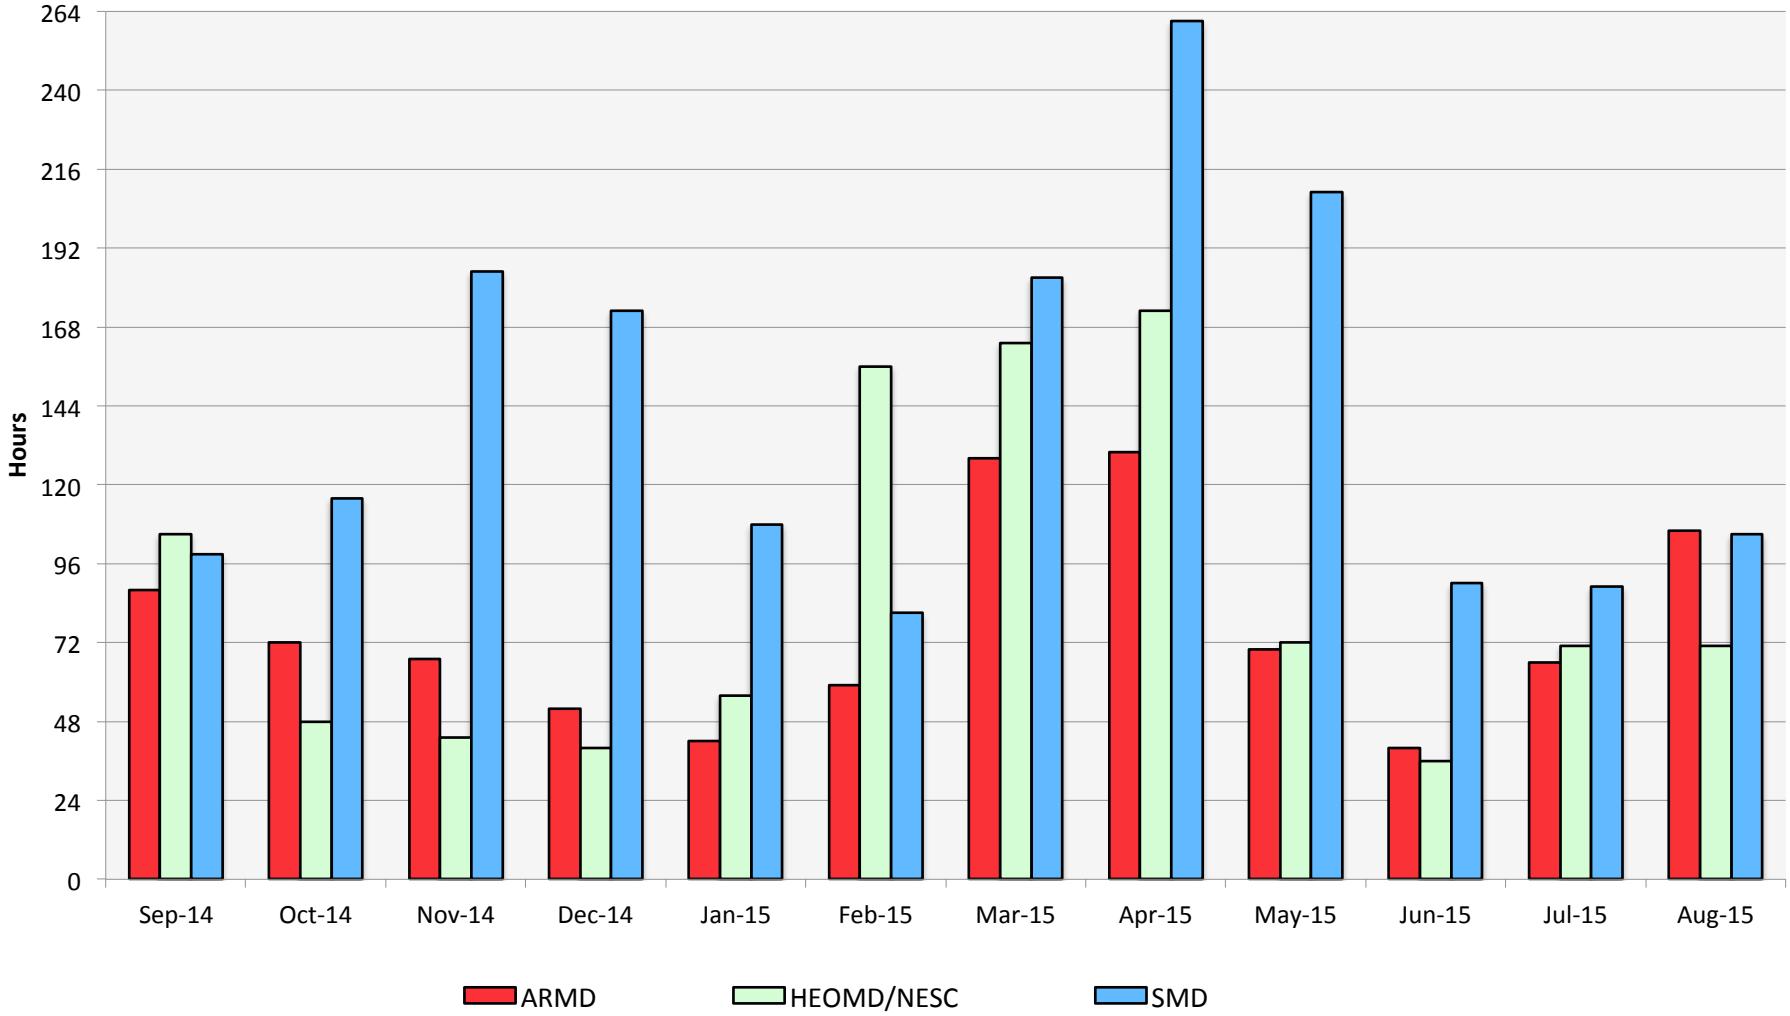

Pleiades:
Average Expansion Factor

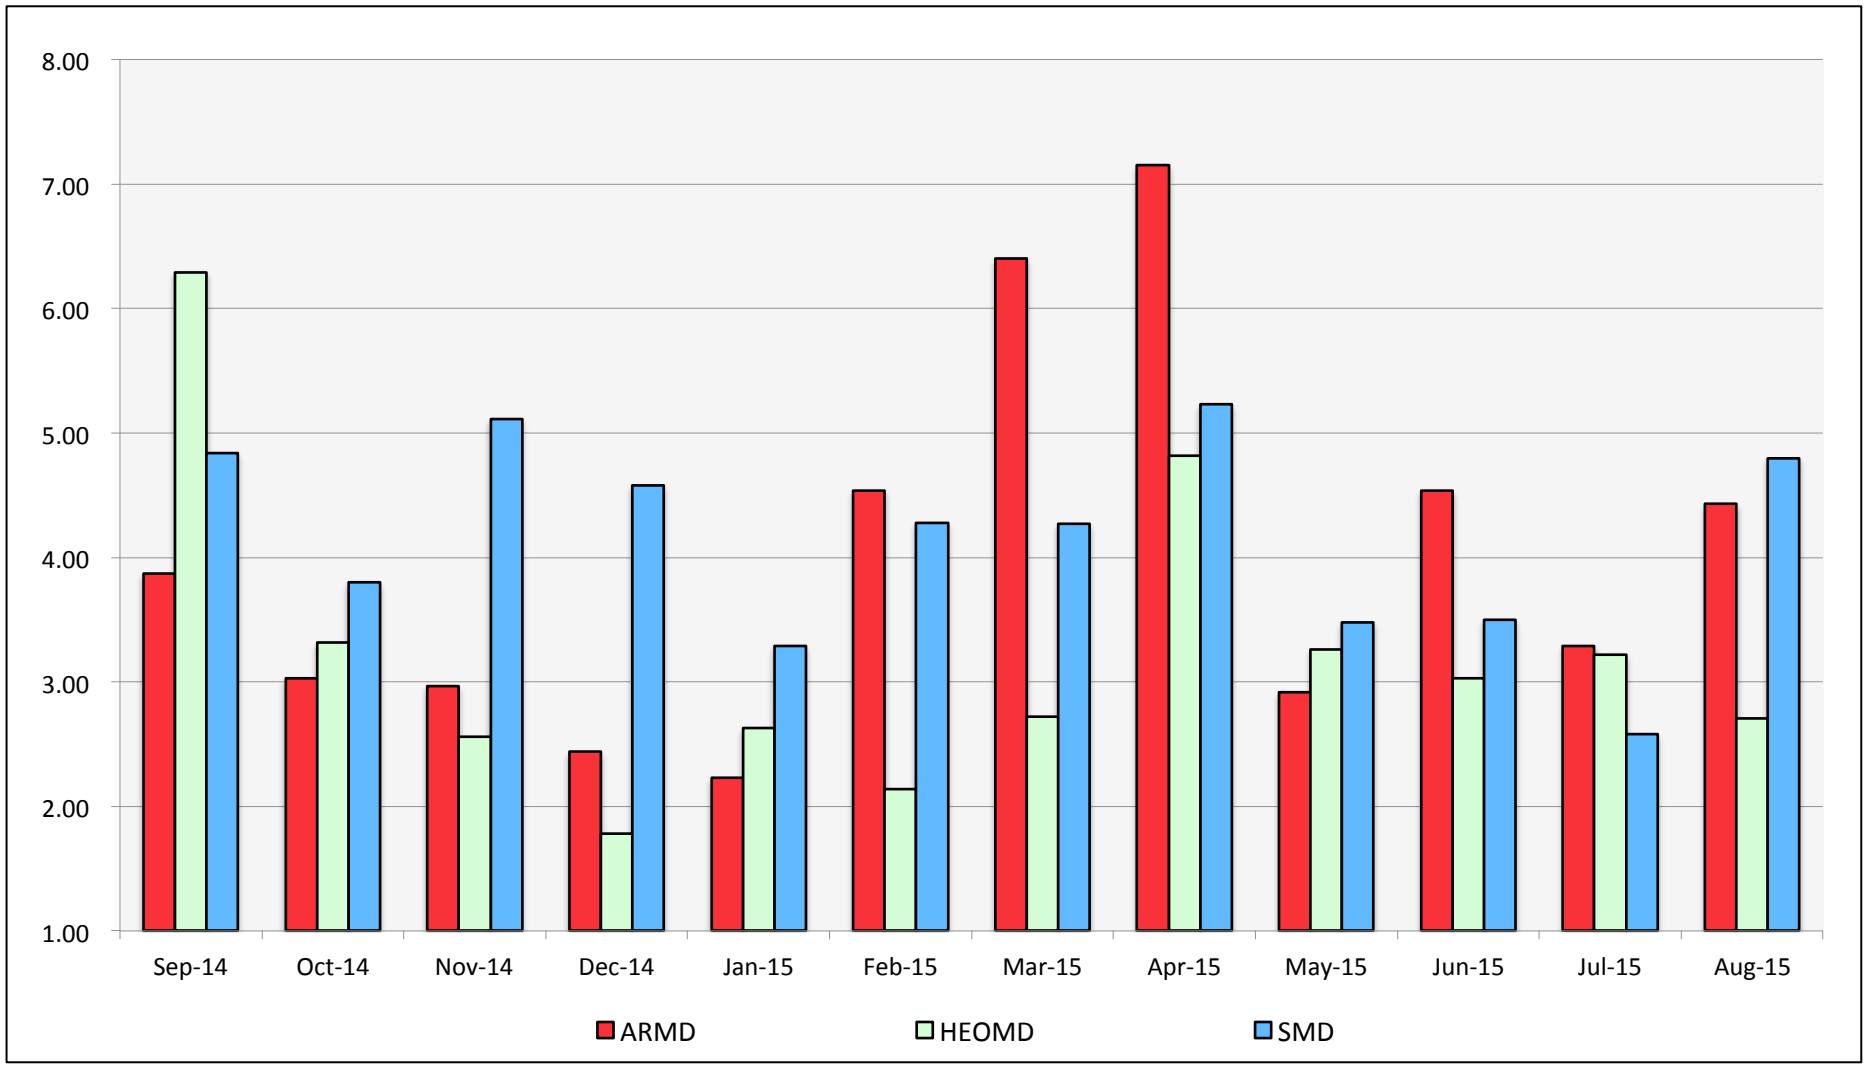

Endeavour:
SBUs Reported, Normalized to 30-Day Month

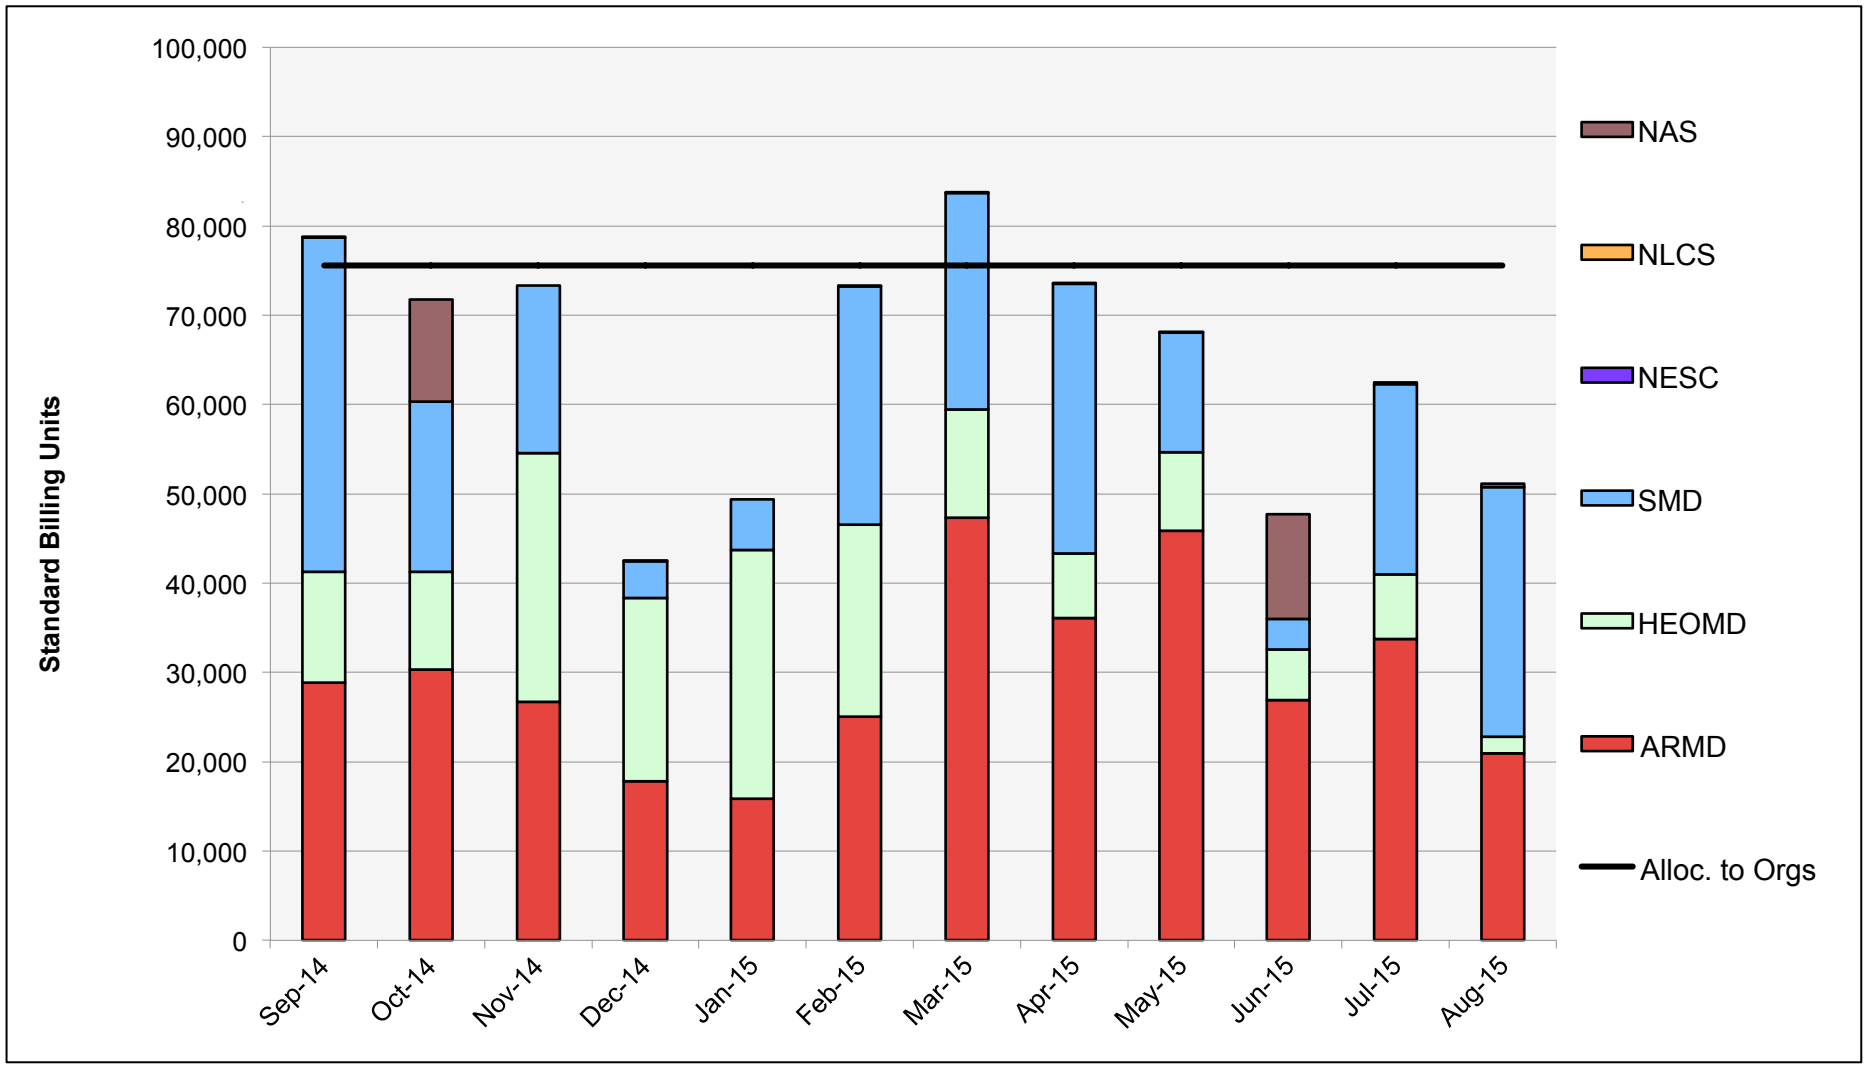

Endeavour:
Monthly Utilization by Job Length

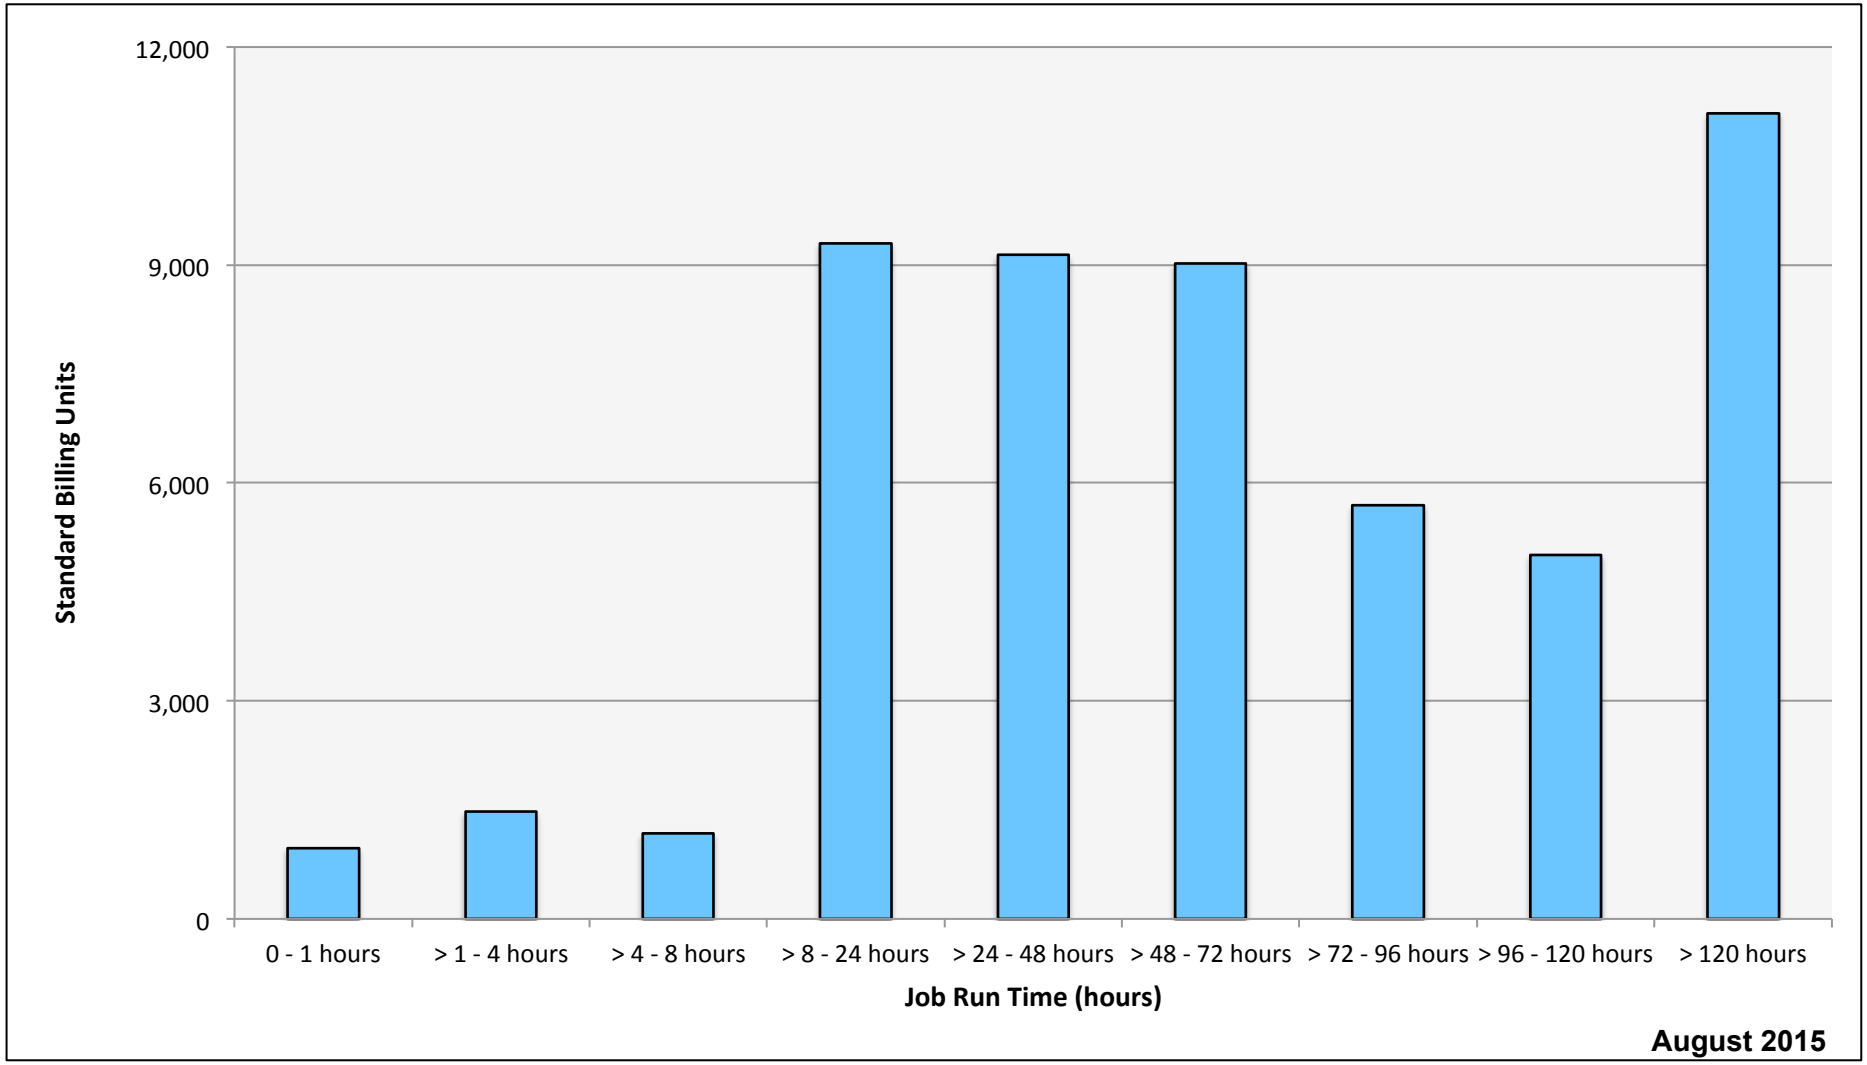

Endeavour:
Monthly Utilization by Size and Mission

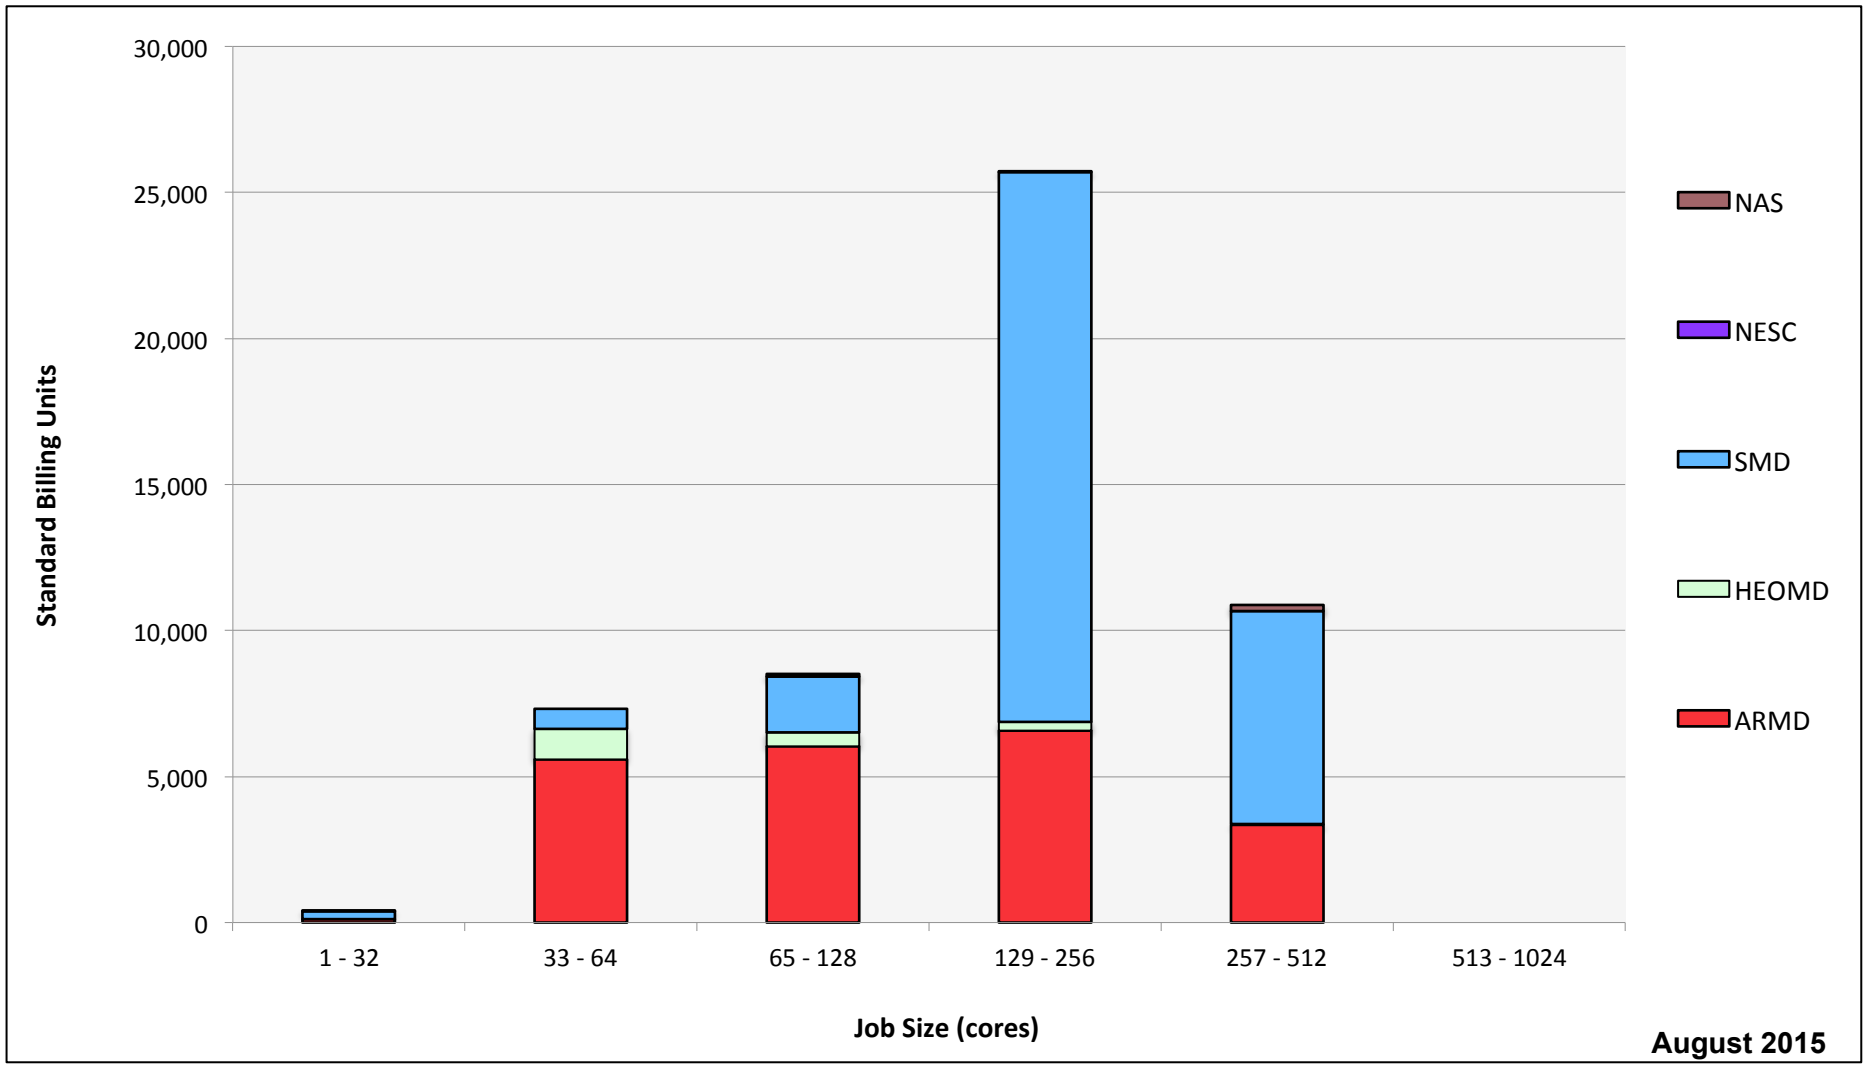

Endeavour:
Monthly Utilization by Size and Length

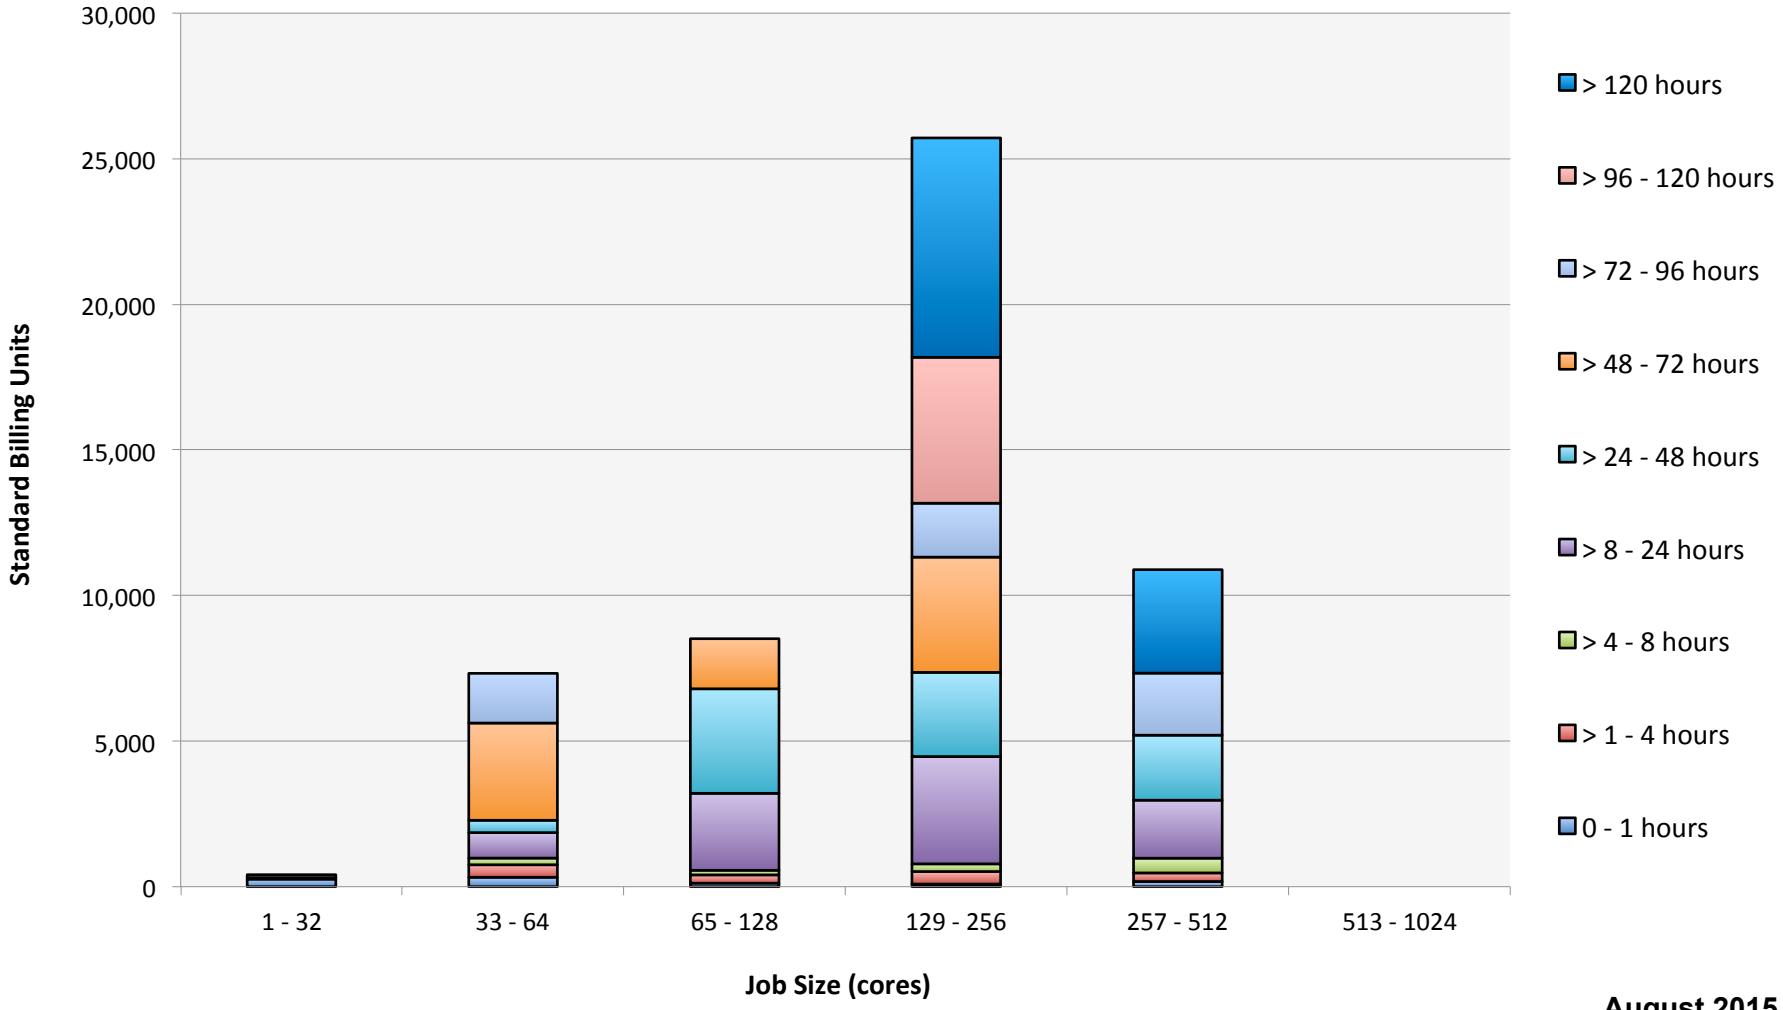

Endeavour:
Average Time to Clear All Jobs

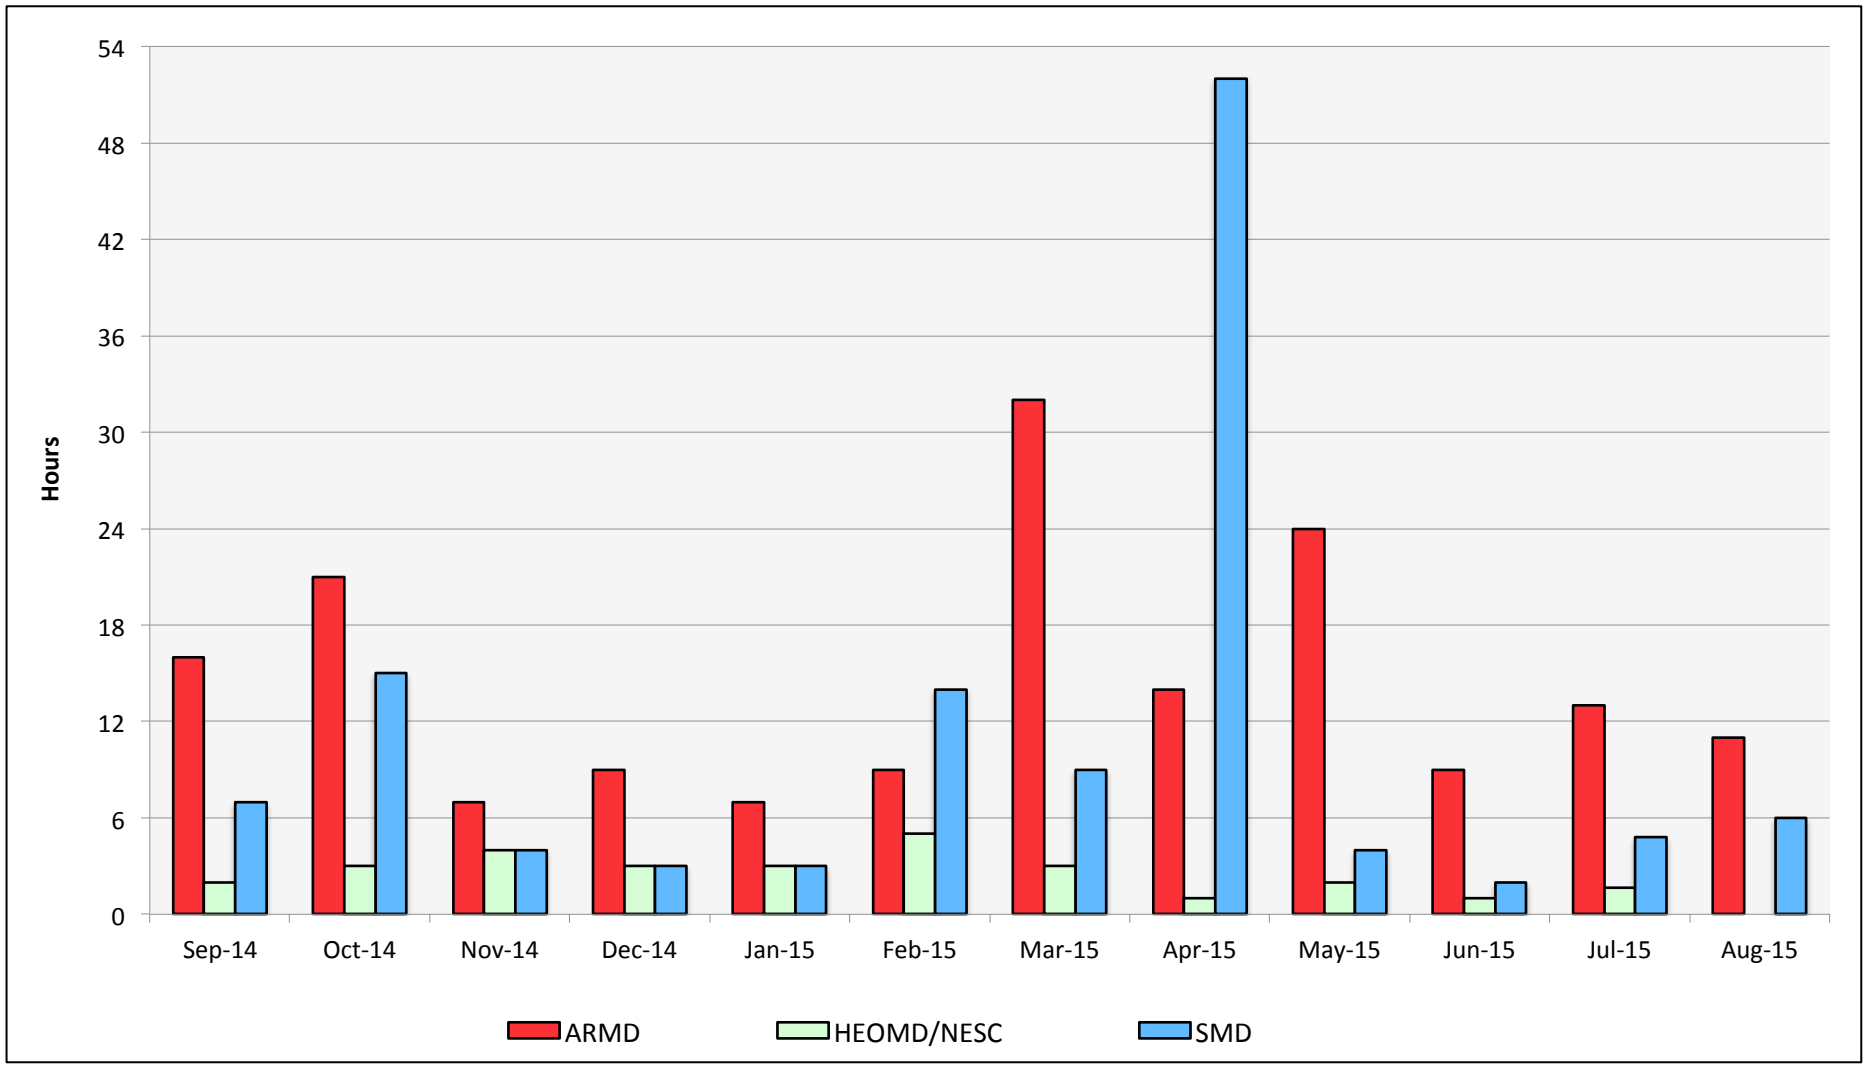

Endeavour:
Average Expansion Factor

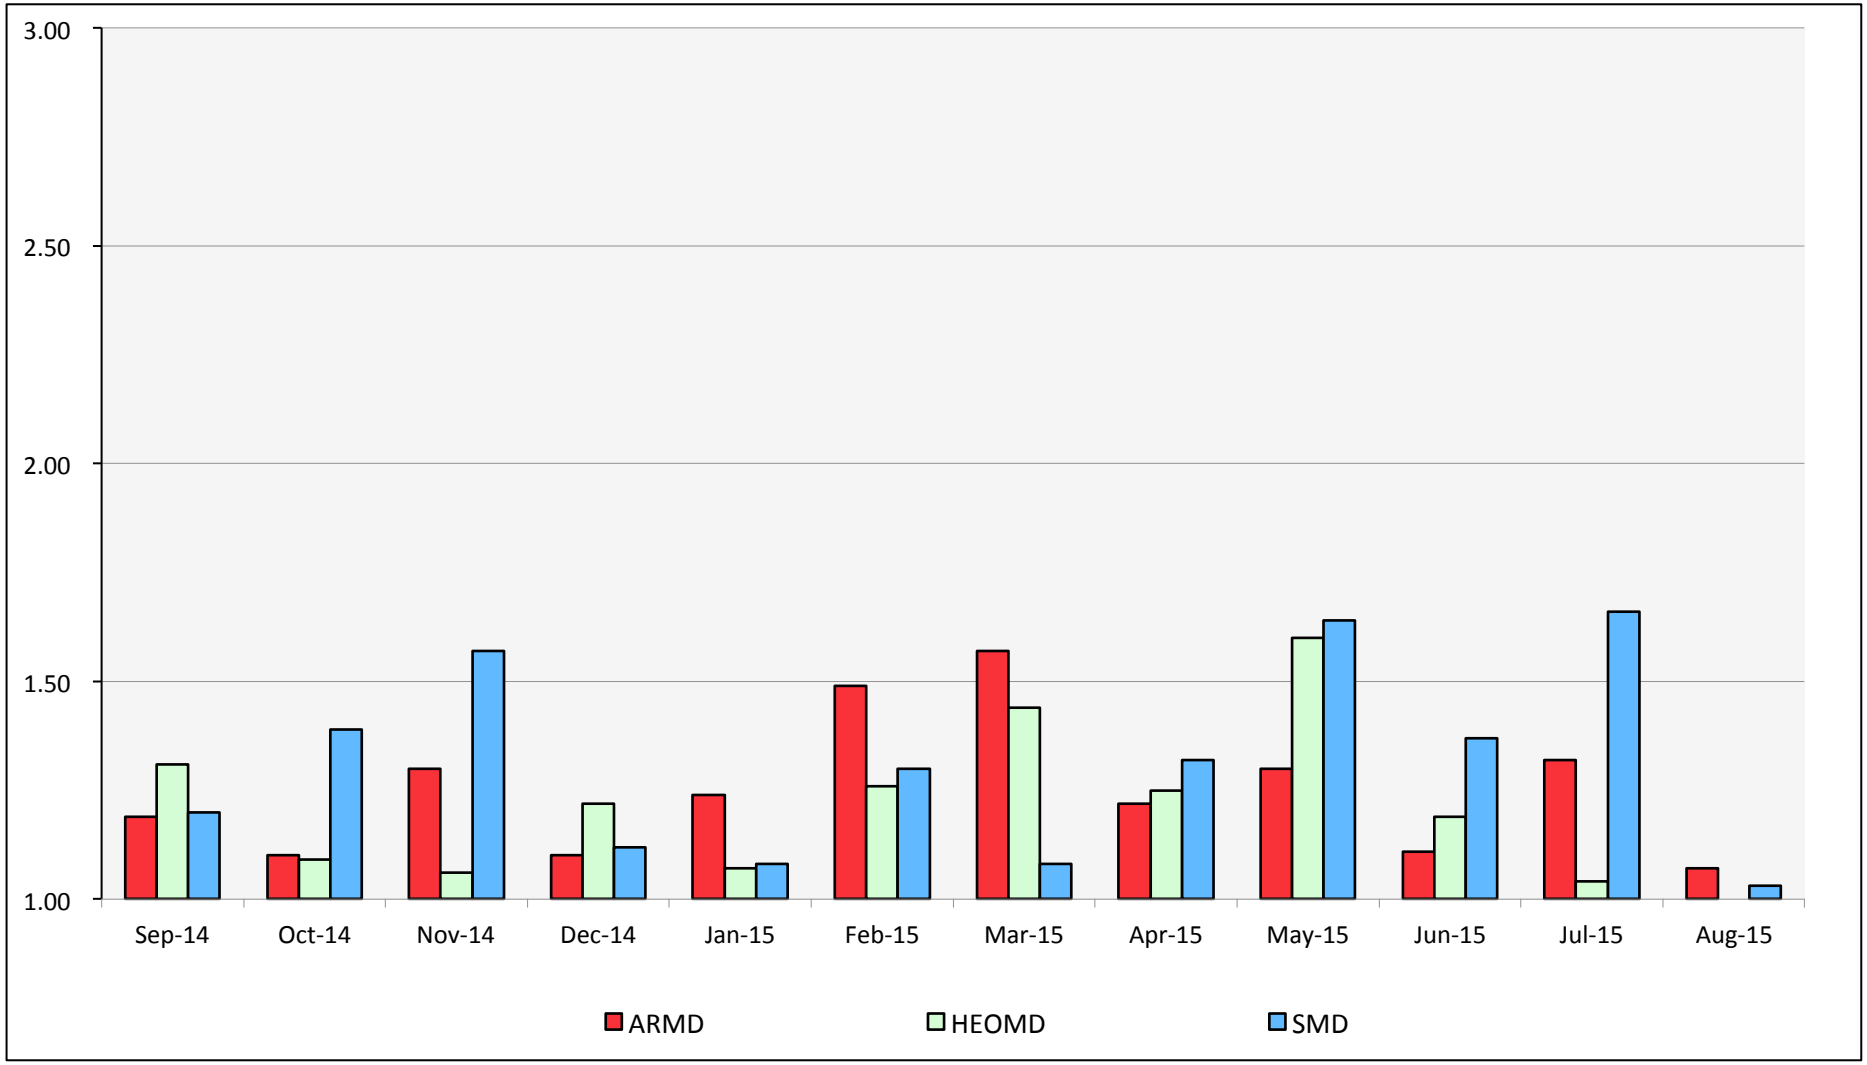

Merope:
SBUs Reported, Normalized to 30-Day Month

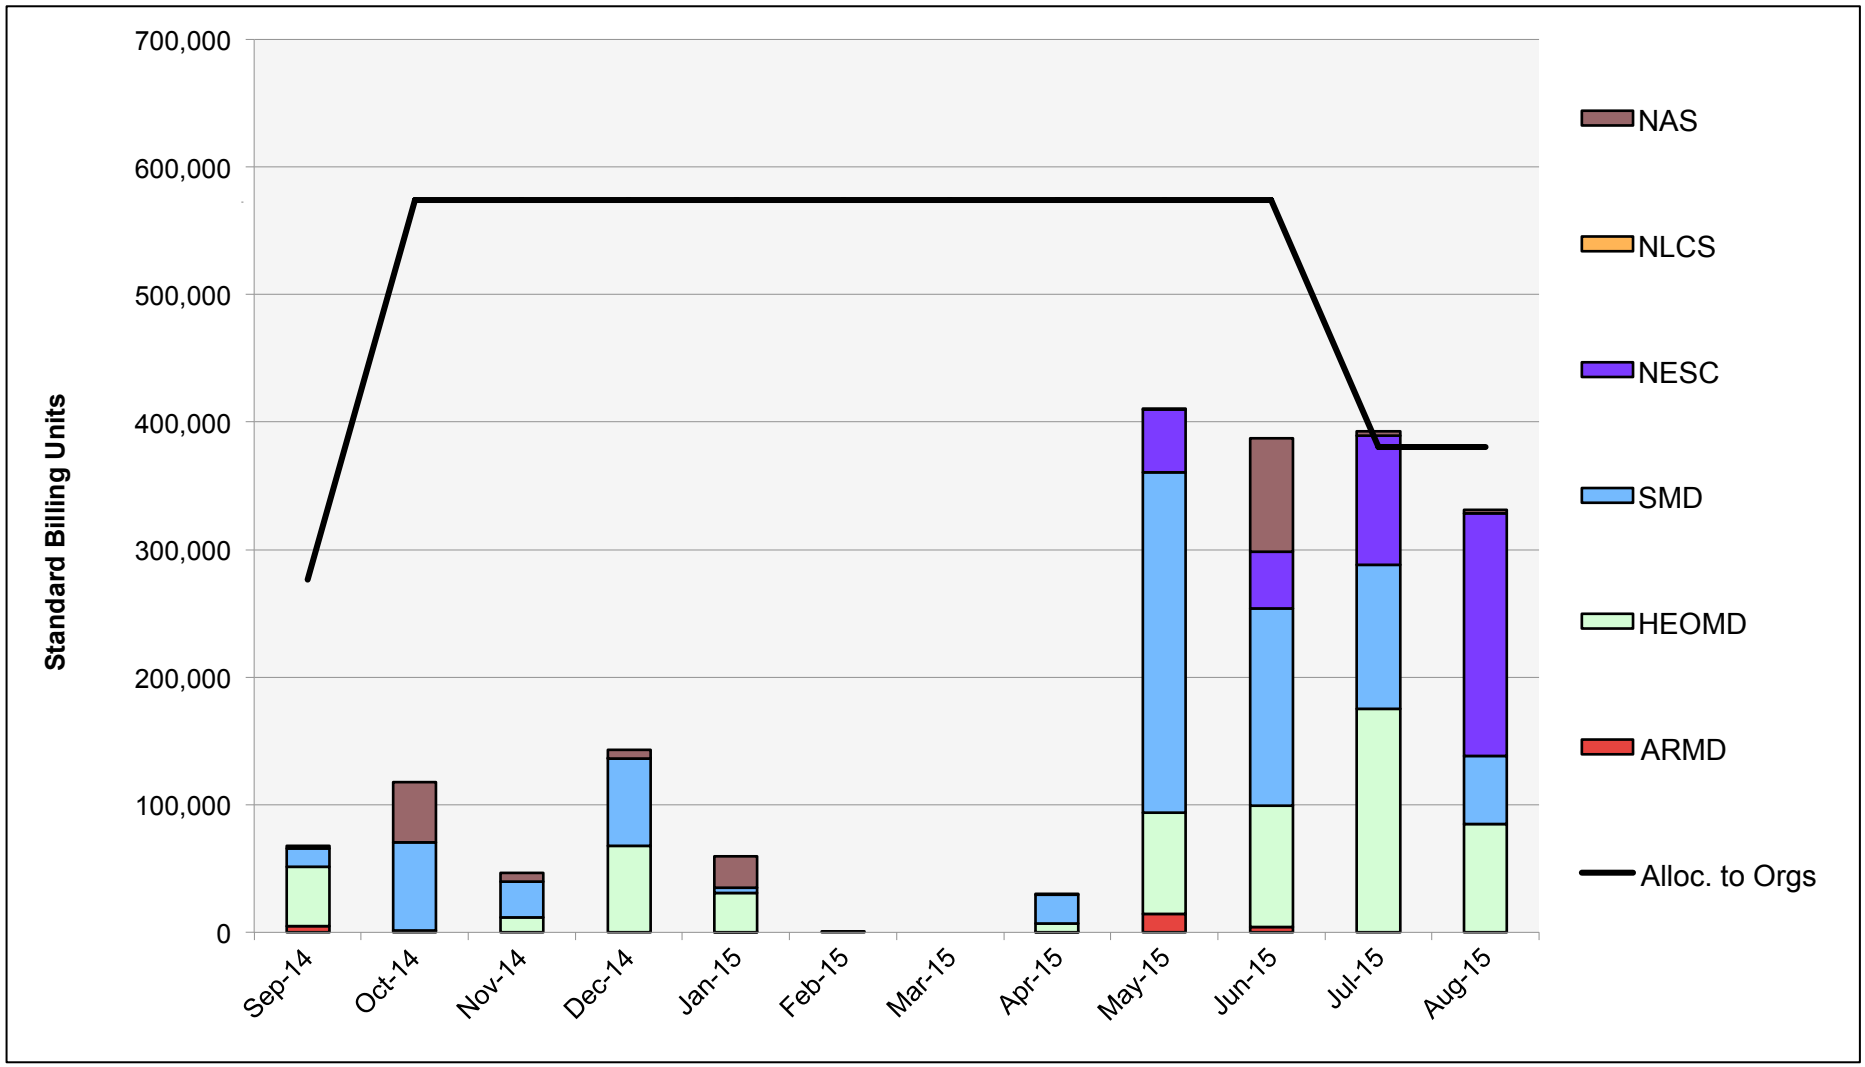

Merope:
Monthly Utilization by Job Length

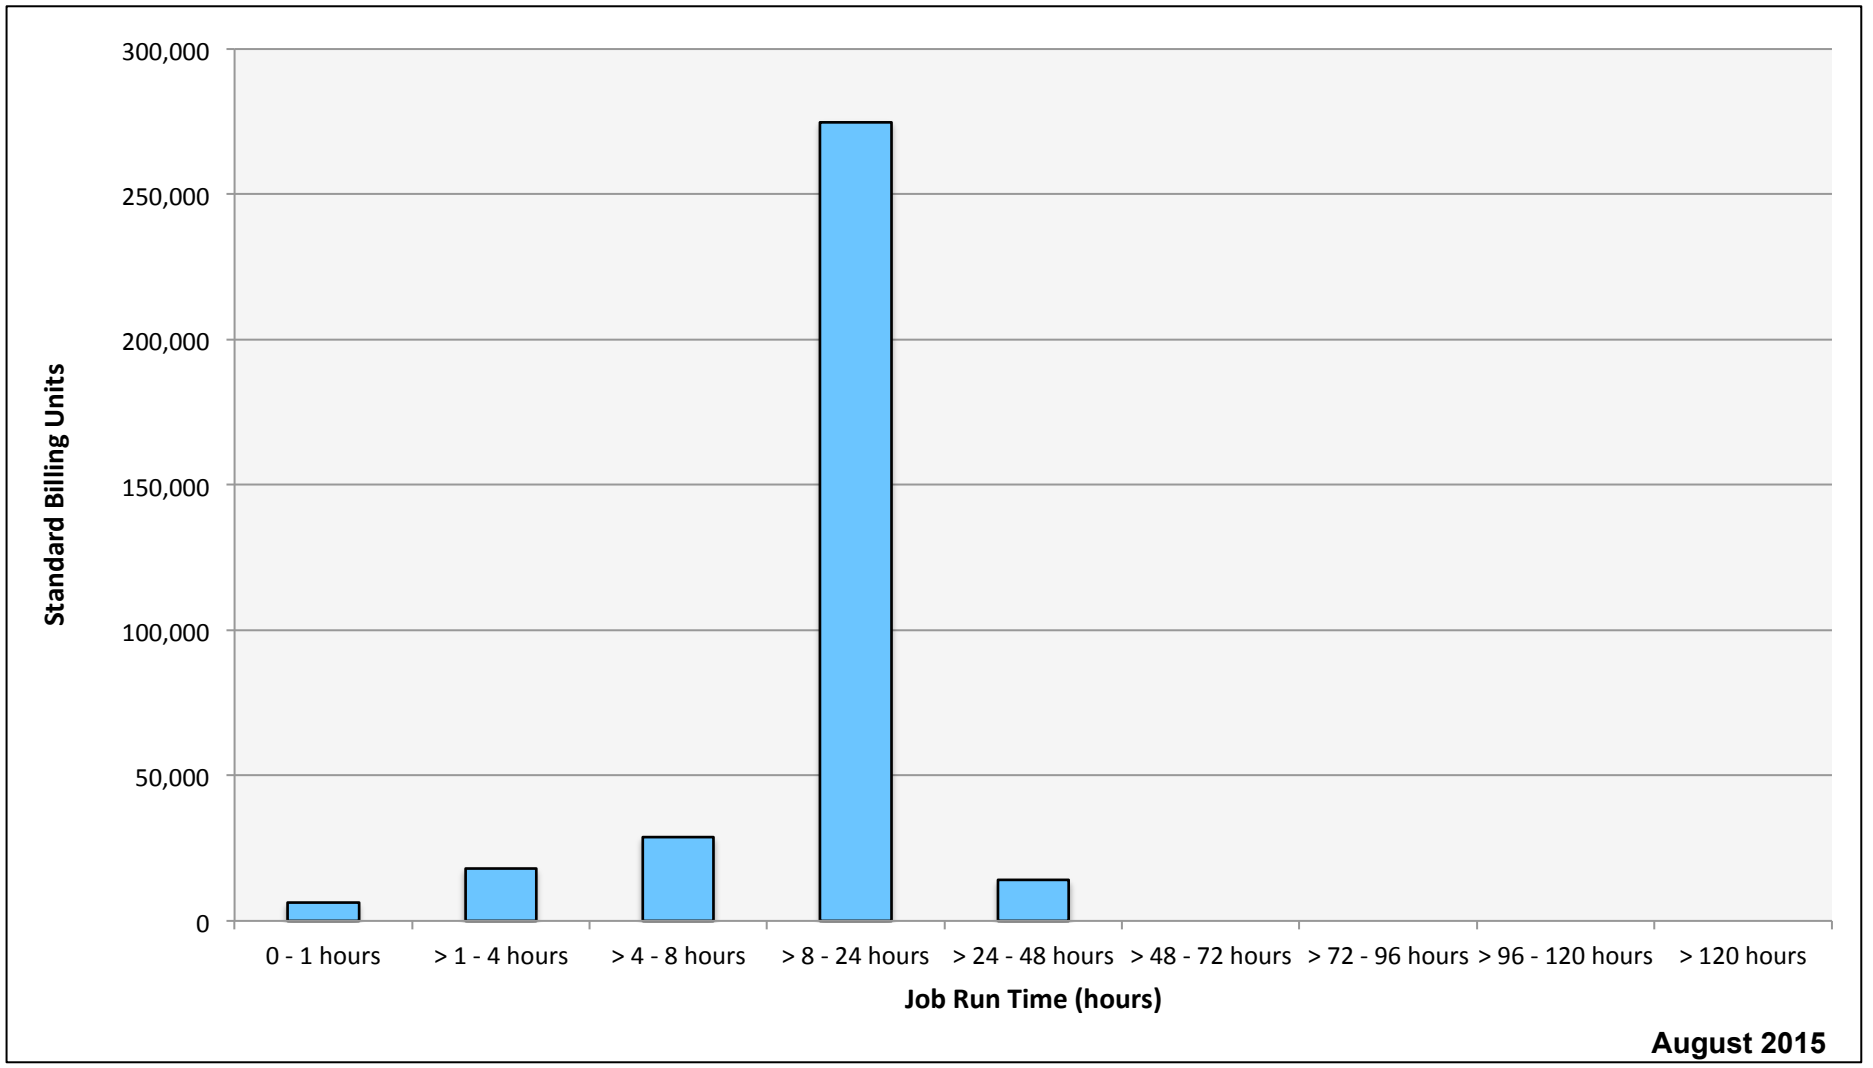

Merope:
Monthly Utilization by Size and Mission

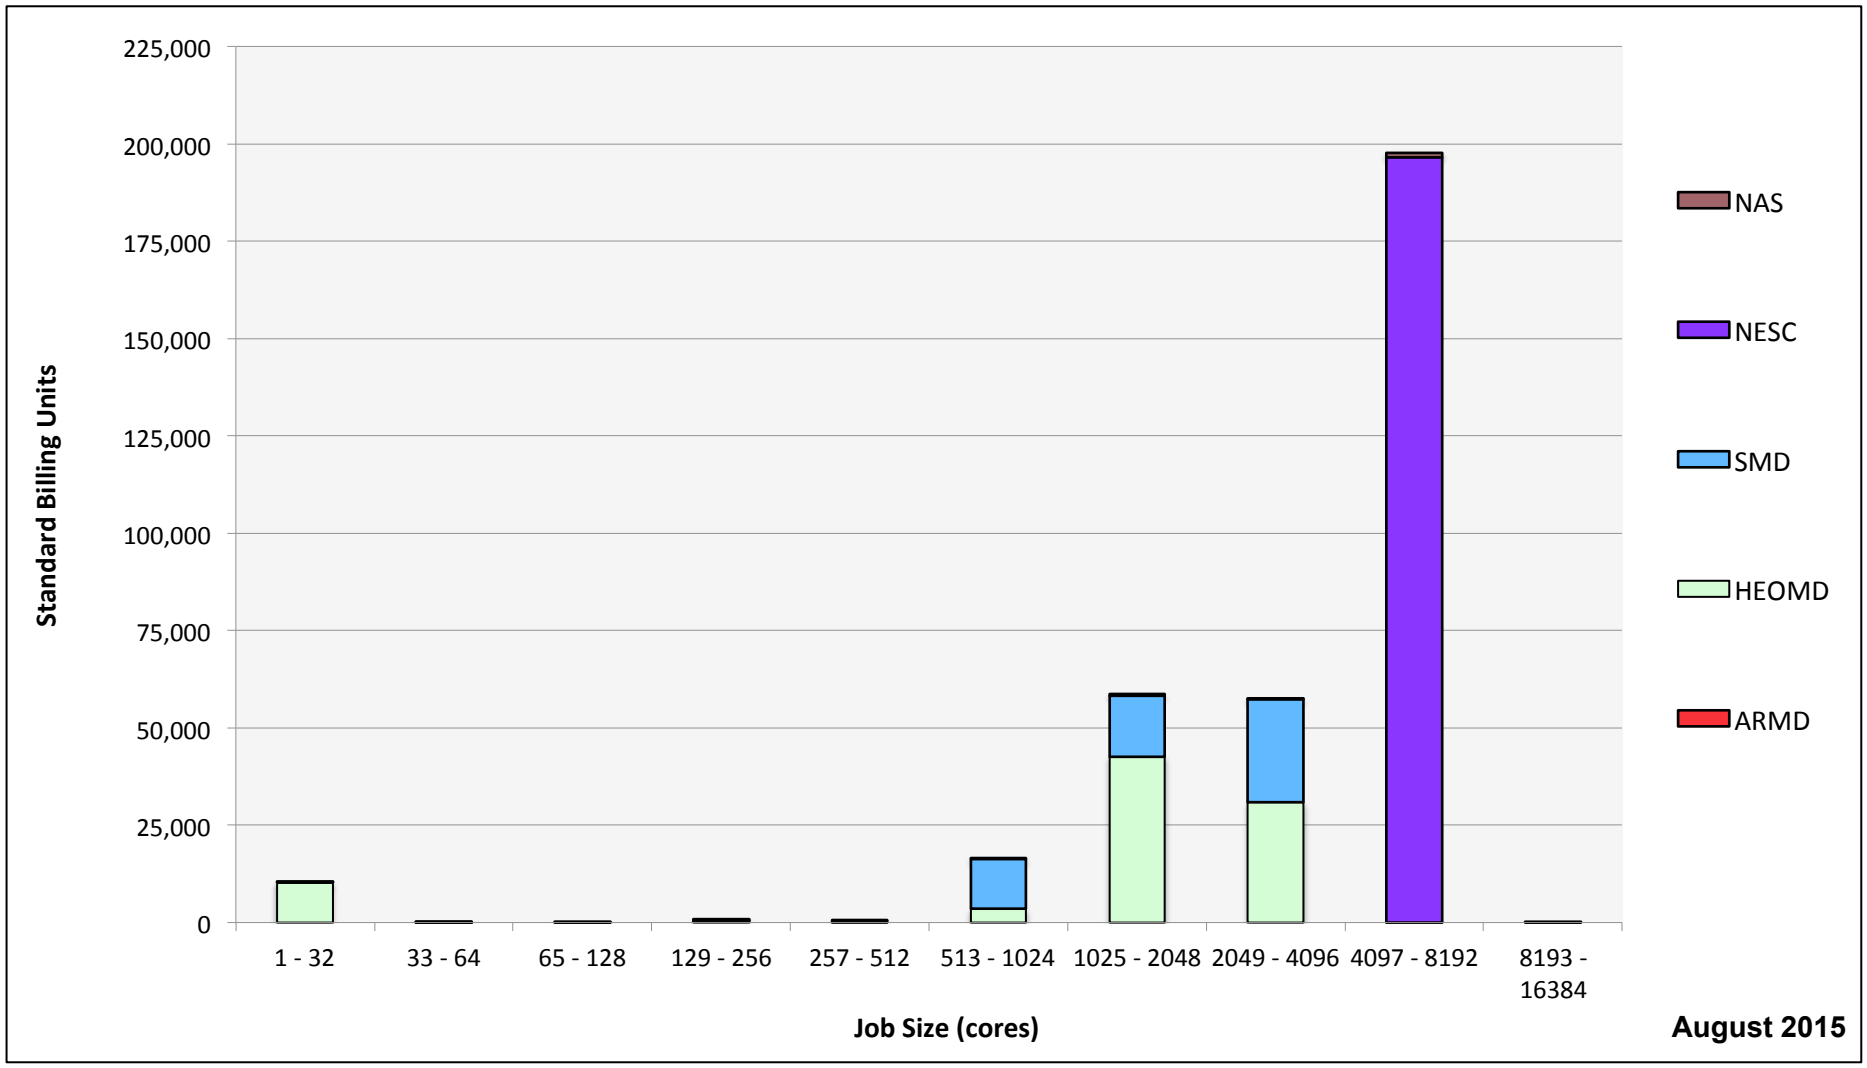

Merope:
Monthly Utilization by Size and Length

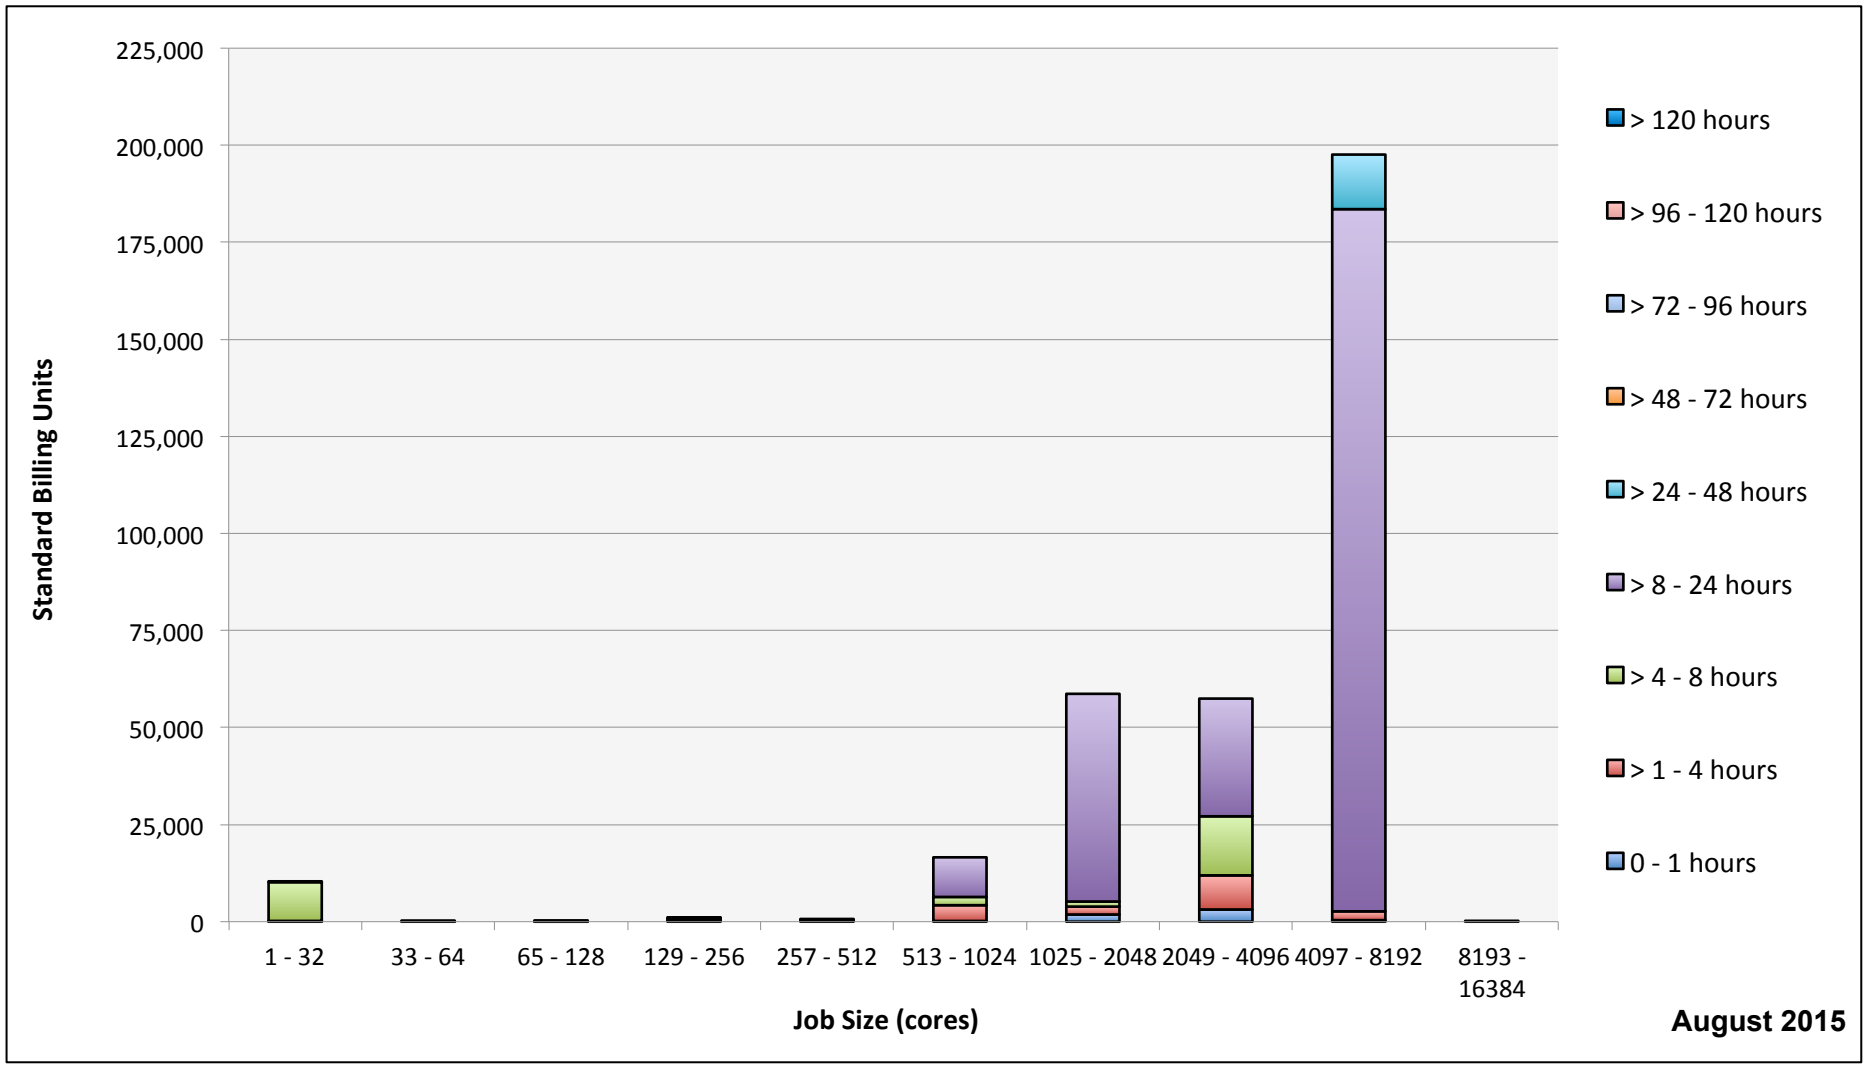

Merope:
Average Time to Clear All Jobs

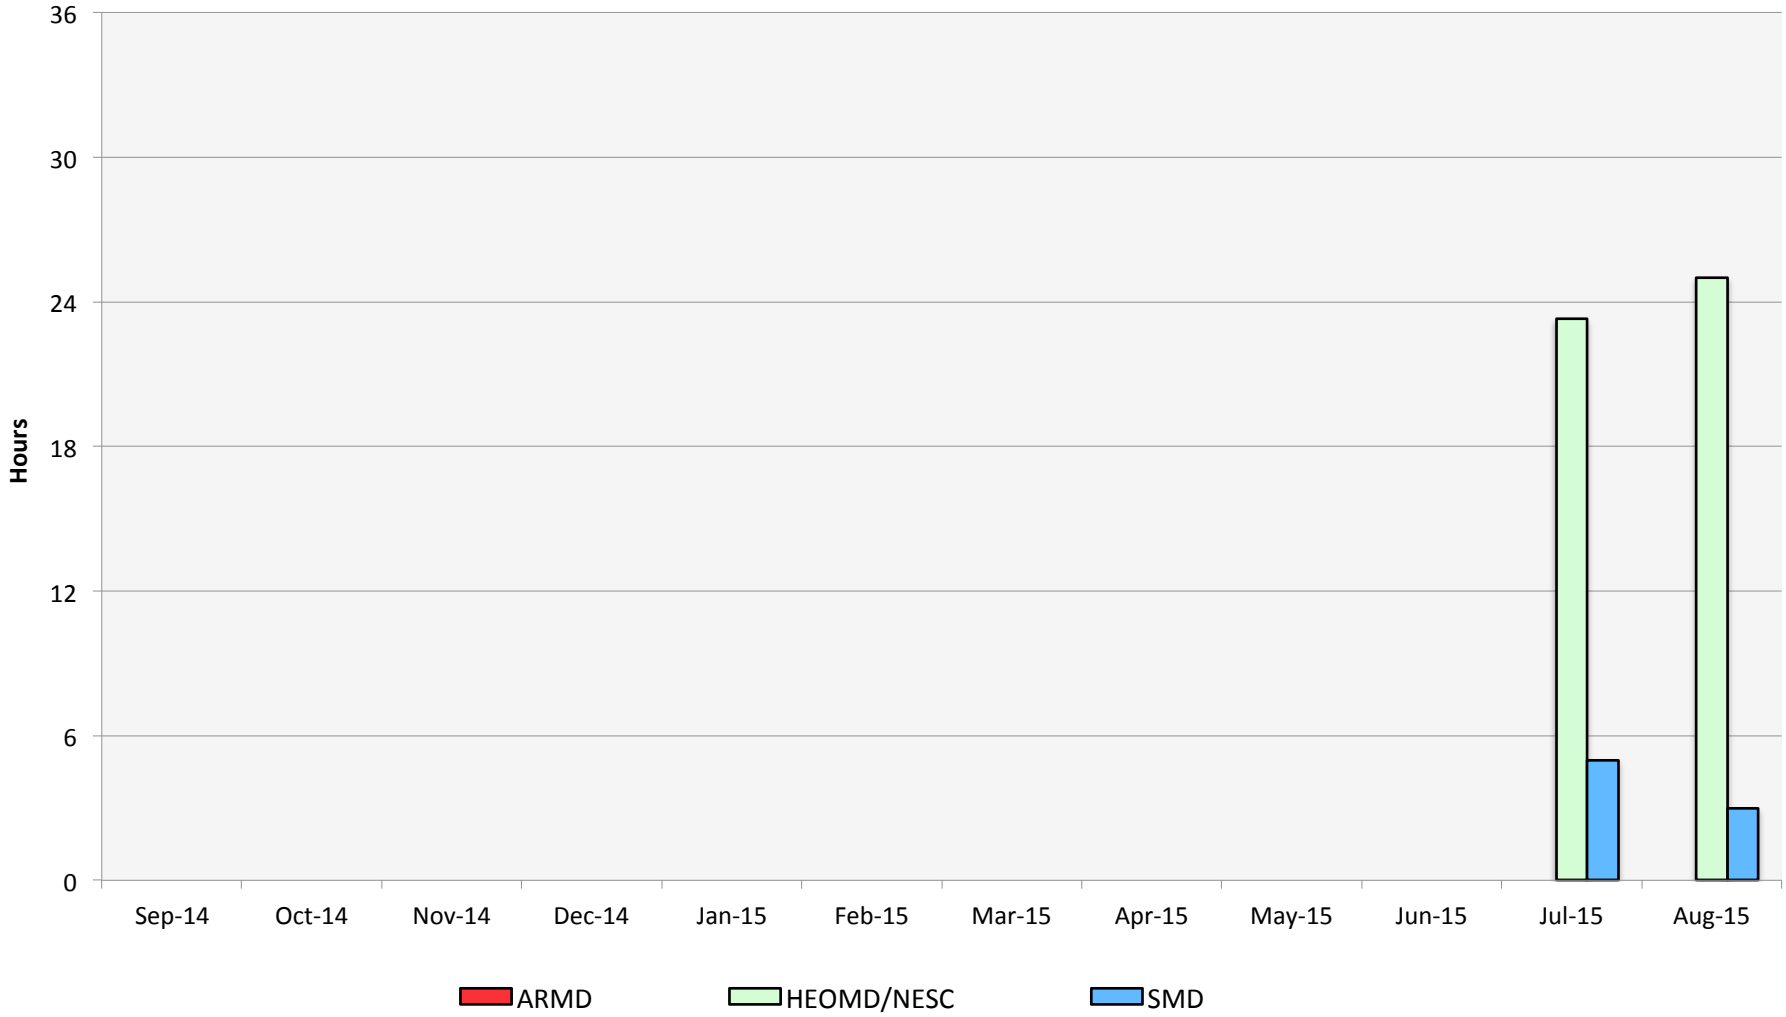

Merope:
Average Expansion Factor